arXiv:1602.03990v2 [stat.ME] 21 Feb 2016

Efficient functional ANOVA through wavelet-domain Markov groves

Li Ma
Duke University

Jacopo Soriano
Duke University

February 23, 2016

Abstract

We introduce a wavelet-domain functional analysis of variance (fANOVA) method based on a Bayesian hierarchical model. The factor effects are modeled through a spike-and-slab mixture at each location-scale combination along with a normal-inverse-Gamma (NIG) conjugate setup for the coefficients and errors. A graphical model called the Markov grove (MG) is designed to jointly model the spike-and-slab statuses at all location-scale combinations, which incorporates the clustering of each factor effect in the wavelet-domain thereby allowing borrowing of strength across location and scale. The posterior of this NIG-MG model is analytically available through a pyramid algorithm of the same computational complexity as Mallat's pyramid algorithm for discrete wavelet transform, i.e., linear in both the number of observations and the number of locations. Posterior probabilities of factor contributions can also be computed through pyramid recursion, and exact samples from the posterior can be drawn without MCMC. We investigate the performance of our method through extensive simulation and show that it outperforms existing wavelet-domain fANOVA methods in a variety of common settings. We apply the method to analyzing the orthosis data.

1 Introduction

This work concerns a common inference task—identifying the contributions from various sources to the variation in functional data, or functional analysis of variance (fANOVA) [18, 20]. Suppose functional observations are measured at a number of given locations. (Throughout this work we use "locations" in a general sense to refer to the points in the index or coordinate space on which the function is observed. For example, in time-series applications, these "locations" are time points.) A simple approach to fANOVA is to carry out ANOVA—e.g., through an F-test—at each location. Doing fANOVA in this location-by-location manner often results in poor performance due to the limited amount of data available at each location as well as the necessary multiple testing correction incurred.

An alternative, often more effective approach, is to first apply a basis transformation to the original observations and then carry out ANOVA under the new basis. Many common bases for functional data analysis can be adopted, including splines, polynomials, and Fourier basis, etc. [18]. With a properly chosen basis, this can allow more effective pooling of information across multiple locations, as well as more sparse representations of the underlying structures (here the functional variations), thereby enhancing the ability to identify such structures. The effectiveness of the different bases depends on the nature of the data; no basis is universally the best for all problems. Here our attention is focused on the wavelet basis transform, and in particular the discrete wavelet transform (DWT) [13].

The DWT is a basis transform for functional data observed on an equi-spaced grid of locations. It has been extensively applied in applications such as signal processing and time-series analysis [17]. The transformed data, in the form of the so-called wavelet coefficients, characterize functional features of different scales (also referred to as frequencies) at different locations. Each wavelet coefficient is associated with one location-scale (also referred to as time-frequency) combination. This wavelet-domain representation of functional data enjoys several desirable properties. First, the wavelet transform has a "whitening effect"

that reduces the correlation in the noise, making the common assumption of independent errors more reasonable than in the original space. A second benefit is that wavelet transforms concentrate "energy"—in an information theoretic sense as measured by entropy or Kullback-Leibler divergence—into a small number of location-scale combinations, and thus "sparsify" the underlying signal thereby making detecting such structures much easier. These properties have motivated the development of numerous wavelet-based regression methods in function estimation. Such methods are particularly effective in comparison to other functional methods when the underlying functions contain local structures [15].

More recently, several authors have proposed methods for fANOVA under the DWT [19, 2, 1, 4, 14]. In [19, 14], the authors treat each location-scale combination individually and carry out an ANOVA test for each, and then identify significant p-values after correcting for multiple testing. In a similar vein but instead of taking the coefficient-by-coefficient testing approach, [1] and [2] consider testing of the joint null hypothesis of nil factor effects over all location-scale combinations together, and constructed tests for this purpose that are minimax optimal. On the other hand, while not directly addressing fANOVA, [15] proposes a Bayesian functional mixed-effects model in the wavelet domain that can be applied to this problem. In particular, a regression model is adopted for the wavelet coefficient at each location-scale combination. ANOVA hypothesis testing can then be naturally handled as a model selection problem using the spike-and-slab prior on the inclusion of the fixed/random effects into the regression model. Inference under the model incurs a heavy computational cost requiring Markov Chain Monte Carlo (MCMC) on a large number of regression models, one for each location-scale combination. Instead, [4] introduces a frequentist approach for the wavelet-domain mixed effects models addressing testing in both random and fixed effects.

A key motivation for the current work is an important common phenomenon in wavelet-domain analysis that has not been amply exploited in the existing wavelet-based fANOVA methods—the location-scale clustering of functional features in the wavelet coefficients. That

is, interesting functional features tend to appear in clusters (like a string of grapes) in the wavelet domain. Such correlation structure has been noted as early as [11], and has been fruitfully exploited by [10] in function estimation. It is easy to imagine, and will be confirmed herein, that the location-scale dependency is prevalent in fANOVA problems as well: when a factor impacts the variance of a wavelet coefficient at one location-scale combination, it typically contributes to the variance at nearby locations and in adjacent scales as well. Thus inference should benefit, in fact substantially, from "borrowing strength" among nearby and/or nested location-scale combinations in identifying the variance components.

The Bayesian wavelet regression framework for fANOVA affords a natural, principled way to incorporating such dependency through designing model space priors—that is, priors on the spike-and-slab indicators that encode whether a factor contributes to the variance at different location-scale combinations. In particular, we present a model space prior in the form of a graphical model consisting of a collection of Markov trees [10], one for each factor, and hence called the Markov grove (MG). The MG prior is highly parsimonious—specified by a small number of hyperparameters, and yet flexible enough to characterize the key dependency pattern in factor effects across adjacent/nested location-scale combinations.

Our new Bayesian hierarchical fANOVA model enjoys several important properties. First, due to the tree structure of the MG prior, when coupled with a normal-inverse-Gamma (NIG) conjugate prior specification on the regression coefficients and error variance, exact Bayesian inference for fANOVA can be achieved efficiently. In particular, we show that the joint posterior of our model has a closed form representation computable using a pyramid algorithm that operationally imitates Mallat's pyramid algorithm for the DWT [13] and achieves the same computational complexity (or simplicity rather), being linear in both the number of functional observations and the number of locations. The closed form posterior allows direct sampling from the posterior using standard Monte Carlo as opposed to MCMC. Furthermore, when testing fANOVA hypotheses, the posterior marginal probability for the

alternative hypotheses (i.e., the presence of factor effects) can also be computed analytically using pyramid recursion without Monte Carlo. This makes our model particularly favorable in large-scale problems such as genomics where fANOVA needs to be completed many times.

The rest of the paper is organized as follows. In Section 2 we present our methodology. First, we provide a brief background on Bayesian wavelet regression in Section 2.1. There we review the NIG conjugate prior and show how to use it in conjunction with the MT model to achieve adaptive shrinkage in the wavelet coefficients in a way that takes into account location-scale dependency. In Section 2.2 we introduce the MG model as a generalization to the MT model, and show how to use it with the NIG specification to form a hierarchical model for wavelet-based fANOVA. We construct a full inference framework for fANOVA under this model consisting of (i) a closed form of the joint posterior computable through a pyramid algorithm, (ii) a recipe for evaluating the posterior marginal alternative probability of each factor effect at each location-scale combination based on another pyramid algorithm, and (iii) a decision rule for calling significant factor effects that properly adjusts for multiple testing. In Section 3 we carry out simulations to evaluate the performance of our method and compare it to a number of wavelet-domain fANOVA methods. We also apply our method to the analysis of the orthosis data. We conclude in Section 4 with brief remarks.

2 Method

2.1 Wavelet regression with normal-inverse-Gamma Markov tree

We start from considering Bayesian modeling of a single functional observation in the wavelet domain. We shall use this simpler problem as a medium to introduce a number of building blocks of our more general wavelet-based fANOVA method—namely, (i) Bayesian adaptive wavelet shrinkage with the spike-and-slab prior, (ii) the normal-inverse-Gamma (NIG) conjugate specification, and (iii) the Markov tree (MT) model. We will show how one can use these three tools in conjunction to carry out adaptive shrinkage in the wavelet domain. Our

approach arises from a recombination of the ideas from [7], [9], [10], [5], and [15].

Suppose we have a single functional observation whose values are attained at T equidistant locations $\boldsymbol{y} = (y_1, \ldots, y_T)$, and

$$\begin{aligned} \boldsymbol{y} &= \boldsymbol{f} + \boldsymbol{\epsilon} \\ \boldsymbol{\epsilon} &\sim \mathrm{N}(\boldsymbol{0}, \Sigma_\epsilon), \end{aligned} \tag{1}$$

where $\Sigma_\epsilon = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_T^2)$. In words, the errors are assumed to be independent across the locations but can be heterogeneous. We wish to recover the unknown function $\boldsymbol{f}$ (or some features of it) from the noisy observation $\boldsymbol{y}$. For simplicity, we assume that $T = 2^{J+1}$ for some integer J. After applying the DWT to $\boldsymbol{y}$ we obtain:

$$\boldsymbol{d} = \boldsymbol{z} + \boldsymbol{u},$$

where $\boldsymbol{d} = \boldsymbol{y}W'$, $\boldsymbol{z} = \boldsymbol{f}W'$ and $\boldsymbol{u} = \boldsymbol{\epsilon}W'$ with W being the orthonormal matrix corresponding to the corresponding wavelet basis. Due to properties of multivariate Gaussians, $\boldsymbol{u}$ is also a multivariate Gaussian with a diagonal covariance matrix.

The elements of $\boldsymbol{d} \in \mathbb{R}^T$ are referred to as the (empirical) wavelet coefficients, and those in $\boldsymbol{z} \in \mathbb{R}^T$ are the wavelet coefficients of its mean function $\boldsymbol{f}$. In particular, one of the elements in each of $\boldsymbol{d}$, $\boldsymbol{z}$, and $\boldsymbol{u}$ is called the *father* (or scaling) coefficient, while the other $T-1$ are the (mother) coefficients. Without loss of generality, in this work we assume that the scaling coefficients are computed at the coarsest level. The elements of $\boldsymbol{d}, \boldsymbol{z}$ and $\boldsymbol{u}$ can be organized into a bifurcating tree structure, with each element in the corresponding vector associated to a node in the tree. We use the pair of indices (j,k), where $j = 0, \ldots, J = \log_2 T - 1$ and $k = 0, \ldots, 2^j - 1$, to represent the kth node in the jth level of this tree. The two children nodes of node (j,k) are indexed by $(j+1, 2k)$ and $(j+1, 2k+1)$. Correspondingly, for $j > 1$, the parent of (j,k) is indexed by $(j-1, \lfloor k/2 \rfloor)$. From now on we shall use *node* and *location-scale combination* interchangeably, and use $\mathcal{T}$ to denote the collection of indices (j,k) corresponding to all nodes in the bifurcating tree. We shall use $d_{j,k}$, $z_{j,k}$, and $u_{j,k}$ to

denote the corresponding mother wavelet coefficients.

The model can be written in a node-specific manner (for notational simplicity, we express the model in terms of the mother coefficients, but the same holds for the father coefficients):

$$d_{j,k} = z_{j,k} + u_{j,k} \quad \text{where} \quad u_{j,k} \sim \mathrm{N}(0, \sigma^2_{j,k}).$$

Bayesian inference on this regression model proceeds by placing priors on $z_{j,k}$ as well as on the hyperparameter $\sigma^2_{j,k}$.

It is well-known that effective inference in the wavelet regression should exploit the underlying sparsity of the wavelet coefficients—that is, many of the coefficients $z_{j,k}$ are (or very close to) zero—due to energy concentration. Hence, data-adaptive shrinkage toward zero is critical for effectively inferring $z_{j,k}$. A very popular Bayesian strategy to achieving this, which has been adopted by several authors, is to place a two-group, or so-called spike-and-slab, mixture prior on $z_{j,k}$ [7, 9, 8, 5, 15]:

$$z_{j,k} \sim (1 - \pi_{j,k}) \cdot \delta(0) + \pi_{j,k} \cdot N(0, \tau_j \sigma^2_{j,k}).$$

In words, with prior probability $\pi_{j,k}$, the wavelet coefficient $z_{j,k}$ is non-zero and its value is generated from a Gaussian distribution. (More generally, the spike does not have to be exactly at 0 but can be a Gaussian with a much smaller variance, to which our method will also apply.) The hyperparameter τ_j is a level-specific dispersion parameter that characterizes the overall level of variability in the wavelet coefficients at level j. Specifically, we consider the following parametric structure as proposed in [3]:

$$\tau_j = 2^{-\alpha j} \tau,$$

for some $\alpha, \tau > 0$. This implies that the wavelet coefficients tend to be smaller as the level j increases. The parameter α controls the smoothness of the functional observations. Larger α corresponds to smoother functions. In particular, [3] discusses guiding principles for selecting α to generate functions of various regularities. A moderate choice that works

well for a variety of common functions, as recommended in [3], is $\alpha = 0.5$. Alternatively, α and τ can both be chosen through an empirical Bayes approach by maximizing the marginal likelihood (see Section 2.3 for details), which as we will illustrate in the numerical examples can often lead to better performance through incorporating additional adaptivity.

The spike-and-slab prior can be written hierarchically with the introduction of a hidden state $S_{j,k} \in \{0, 1\}$ [8], with $S_k = 1$ indicating that the empirical wavelet coefficient $d_{j,k}$ contains a signal $z_{j,k} \neq 0$, and when $S_k = 0$ otherwise. Formally,

$$z_{j,k}|S_{j,k} \sim \mathrm{N}(0, S_{j,k} \cdot \tau_j \sigma_{j,k}^2) \tag{2}$$

The error variance $\sigma_{j,k}^2$ is typically unknown, and it can be inferred from the data. In many applications of wavelet regression, the error variance is assumed to be homogeneous, i.e., $\sigma_{j,k}^2 \equiv \sigma_0^2$ for all j and k. It has been noted that homogeneous error variance is often unrealistic [15]. Thus we allow $\sigma_{j,k}^2$ to be heterogeneous and adopt a hyperprior on them:

$$\sigma_{j,k}^2 \overset{\text{ind}}{\sim} \text{Inv-Gamma}(\nu + 1, \nu\sigma_0^2). \tag{3}$$

The inverse-Gamma prior maintains conjugacy to the Gaussian model, and consequently the marginal likelihood can be evaluated analytically. This hierarchical specification includes the homogeneous variance as a special case because as $\nu \to \infty$, $\sigma_{j,k}^2 \xrightarrow{p} \sigma_0^2$.

Donoho and Johnstone [11] noted a prevalent phenomenon in many applications of inference in wavelet spaces: the "signals"—the wavelet coefficients that are large in magnitude—often show up in clusters in the location-scale tree. This phenomenon, for instance, can be clearly seen in the four test functions presented in [11] (see Figure 1). In particular, when the coefficient $z_{j,k}$ deviates far away from zero, the coefficients of the two children in the bifurcating location-scale tree, namely $z_{j+1,2k}$ and $z_{j+1,2k+1}$ tend to be away from zero as well. Such location-scale clustering is particularly strong for functions with sharp boundaries and abrupt changes such as blocks, bumps, and doppler. Crouse *et al* [10] pointed out that the clustering pattern in the wavelet coefficients can be directly exploited to improve

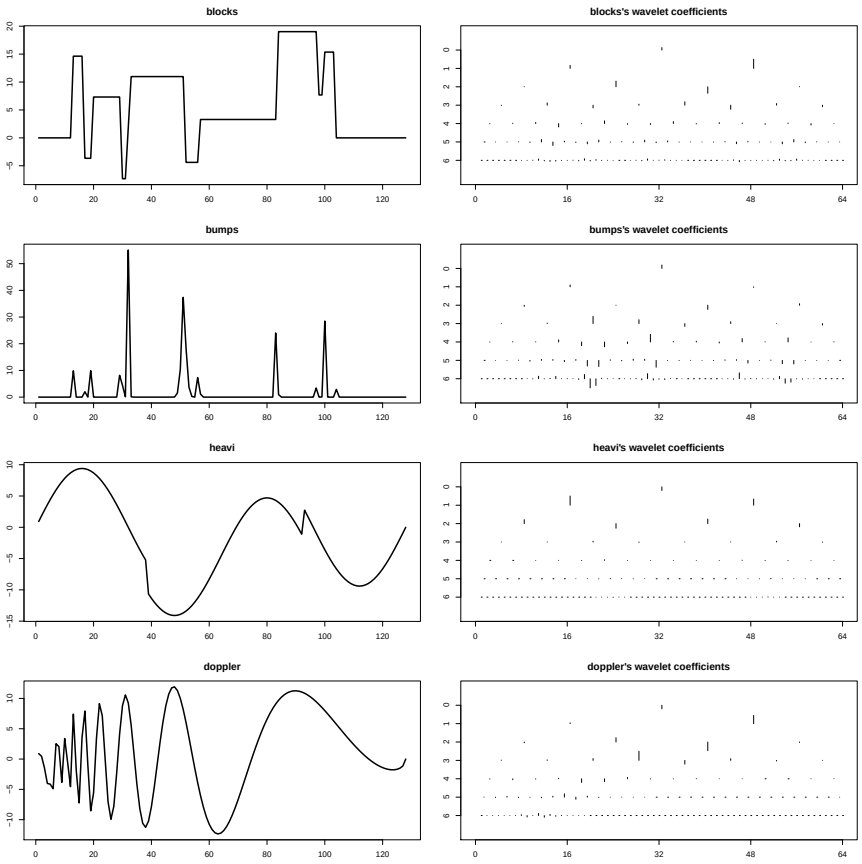

Figure 1: The four test functions from [11] and the associated mother wavelet coefficients. The coefficients for the test functions *blocks*, *bumps* and *doppler* display strong patterns of location-scale clustering.

adaptive shrinkage, and proposed a graphical modeling strategy to induce such spatial-scale dependency by jointly modeling the latent states $S_{j,k}$ using a Markov process, resulting in a hidden Markov model evolving on the location-scale tree, called the *Markov tree* (MT).

Under the MT for $\boldsymbol{S} = \{S_{j,k} : (j,k) \in \mathcal{T}\}$, the shrinkage state of node (j,k) depends on that of its parent through a Markov transition:

$$\Pr(S_{j,k} = s' | S_{j-1,\lfloor k/2 \rfloor} = s) = \rho_{j,k}(s, s'), \tag{4}$$

where $\rho_{j,k}(s,s')$ for $s, s' \in \{0,1\}$ are called the *state transition probabilities*, and they can be organized into a 2×2 transition matrix, $\boldsymbol{\rho}_{j,k}$ for each node (j,k).

A simple and flexible two-hyperparameter specification of these transition matrices is:

$$\boldsymbol{\rho}_{j,k} = \begin{bmatrix} \rho_{j,k}(0,0) & \rho_{j,k}(0,1) \\ \rho_{j,k}(1,0) & \rho_{j,k}(1,1) \end{bmatrix} = \begin{bmatrix} \max\{1-\eta 2^{-j}, 0\} & \min\{\eta 2^{-j}, 1\} \\ 1-\gamma & \gamma \end{bmatrix}, \tag{5}$$

where $0 < \gamma < 1$ and $\eta > 0$. The parameter γ induces the spatial-scale dependency of the wavelet signal. Larger γ values correspond to stronger correlation or clustering in the large wavelet coefficients. On the other hand, the parameter η controls how likely it is to have a "signal", i.e., non-zero wavelet coefficient in each level. The exponential decaying factor 2^{-j} counters exactly the exponential increase in the expected number of wavelet coefficients in higher resolution, and keeps the prior expected number of *de novo* signals (in the sense that a node contains a signal but its parent does not) in each resolution fixed at η.

There are two strategies to choosing the hyperparameters (η, γ). One is to elicit them based on some criteria for multiplicity adjustment and the other is to choose them using an empirical Bayes approach by maximizing the marginal likelihood. We shall discuss both strategies in Section 2.3.

Because the root node $(0,0)$ does not have a parent, the initial state of the process, $S_{0,0}$, is specified by a set of initial state probabilities $\boldsymbol{\rho}_{0,0} = (\rho_{0,0}(0), \rho_{0,0}(1))$ such that:

$$\Pr(S_{0,0} = s) = \rho_{0,0}(s) \quad \text{for } s \in \{0,1\}. \tag{6}$$

Combining the MT model [10] on $\boldsymbol{S}$ and the NIG hierarchical setup [9, 8], Eqs. (2), (3), (4), and (6) together give a new hierarchical model for the wavelet coefficients, which we shall refer to as the *normal inverse-Gamma Markov tree*, or NIG-MT.

We next show how to do inference under the NIG-MT model. In particular, we show that the tree nature of the MT combined with the normal inverse-Gamma setup results in full conjugacy of the NIG-MT: the joint posterior on $\{S_{j,k}, z_{j,k}, \sigma^2_{j,k} : (j,k) \in \mathcal{T}\}$ is still an NIG-MT whose parameters can be computed analytically, and can be sampled from directly.

To this end, let us consider a more general case with $n(\geq 1)$ i.i.d. functional observations $\boldsymbol{y}^{(1)}, \boldsymbol{y}^{(2)}, \ldots, \boldsymbol{y}^{(n)}$ from model (1). From now on, we shall use the superscript "(i)" to indicate the terms corresponding to the ith observation. The node-specific model after DWT becomes:

$$d_{j,k}^{(i)} = z_{j,k} + u_{j,k}^{(i)} \quad \text{where} \quad u_{j,k}^{(i)} \sim N(0, \sigma_{j,k}^2).$$

Our interest lies in finding the posterior distribution on $\{S_{j,k}, z_{j,k}, \sigma_{j,k}^2 : (j,k) \in \mathcal{T}\}$ given the observed data. Let $m_{j,k}(s)$ be the marginal likelihood for the node-specific model on (j,k) given that $S_{j,k} = s \in \{0,1\}$:

$$m_{j,k}(s) = \int p(\boldsymbol{d}^{(1)}, \ldots, \boldsymbol{d}^{(n)} | S_{j,k} = s, z_{j,k}, \sigma_{j,k}^2)\, \pi(z_{j,k}, \sigma_{j,k}^2)\, dz_{j,k} d\sigma_{j,k}^2.$$

From the normal-inverse-Gamma conjugacy, the marginal likelihood is in closed form:

$$m_{j,k}(s) = \frac{(\nu\sigma_0^2)^{\nu+1}\Gamma(\nu + n/2 + 1)}{(2\pi)^{n/2}\Gamma(\nu+1)} \cdot \left[\frac{\tau_j^{-1}}{n+\tau_j^{-1}}\right]^{s/2} \cdot \left[\nu\sigma_0^2 + \frac{1}{2}\left(\sum_i (d_{j,k}^{(i)})^2 - s \cdot \frac{(n\bar{d}_{j,k})^2}{n+\tau_j^{-1}}\right)\right]^{-\nu-n/2-1}$$

where $\bar{d}_{j,k} = \sum_i d_{j,k}^{(i)}/n$.

The following theorem shows that the NIG-MT model is completely *conjugate* in the sense that the joint posterior is still an NIG-MT. Moreover, the posterior hyperparamters are available analytically through a recursive algorithm operationally similar to the pyramid algorithm for DWT [13]. From now on, we shall use $\mathcal{D}$ to represent the totality of data.

Theorem 1. The joint posterior on $\{S_{j,k}, z_{j,k}, \sigma_{j,k}^2 : (j,k) \in \mathcal{T}\}$ is still an NIG-MT as follows:

- The posterior of the hidden states $S_{j,k}$ is still a MT:

  1. State transition probabilities:

  $$\Pr(S_{j,k} = s' | S_{j-1,\lfloor k/2 \rfloor} = s, \mathcal{D}) = \rho_{j,k}(s,s') \frac{\phi_{j,k}(s')}{\xi_{j,k}(s)},$$

  for $s, s' \in \{0,1\}$ and $j = 1, 2, \ldots, J$;

2. Initial state probabilities:

$$\Pr(S_{0,0} = s|\mathcal{D}) = \rho_{0,0}(s)\frac{\phi_{0,0}(s)}{\xi_{0,0}(0)} \quad \text{for } s \in \{0,1\}.$$

- The posterior of the variances $\sigma^2_{j,k}$ given $S_{j,k}$ is:

$$[\sigma^2_{j,k}|S_{j,k},\mathcal{D}] \sim \text{Inv-Gamma}\left(\nu + 1 + \frac{n}{2}, \nu\sigma_0^2 + \frac{1}{2}\left(\sum_i (d^{(i)}_{j,k})^2 - \frac{S_{j,k}\cdot(n\bar{d}_{j,k})^2}{n+\tau_j^{-1}}\right)\right).$$

- The posterior of $z_{j,k}$ given $S_{j,k}$ and $\sigma^2_{j,k}$ is:

$$[z_{j,k}|\sigma^2_{j,k}, S_{j,k},\mathcal{D}] \sim \text{N}\left(\frac{S_{j,k}\cdot n\bar{d}_{j,k}}{n+\tau_j^{-1}}, \frac{S_{j,k}\cdot\sigma^2_{j,k}}{n+\tau_j^{-1}}\right).$$

The mappings $\phi_{j,k}$ and $\xi_{j,k} : \{0,1\} \mapsto [0,+\infty)$ are defined recursively in j and can be computed through a bottom-up pyramid algorithm as follows:

$$\phi_{j,k}(s) = \begin{cases} m_{j,k}(s)\cdot\xi_{j+1,2k}(s)\cdot\xi_{j+1,2k+1}(s) & \text{for } j = 0,1,2\ldots,J-1 \\ m_{j,k}(s) & \text{for } j = J, \end{cases}$$

$$\xi_{j,k}(s) = \begin{cases} \sum_{s'\in\{0,1\}}\rho_{j,k}(s,s')\cdot\phi_{j,k}(s') & \text{for } j = 1,2,\ldots,J \\ \sum_{s'\in\{0,1\}}\rho_{0,0}(s')\cdot\phi_{j,k}(s') & \text{for } j = 0. \end{cases}$$

Remark I: The recursive computation of the mappings $\phi_{j,k}$ and $\xi_{j,k}$ is operationally analogous to Mallat's pyramid algorithm [13] for carrying out the DWT. In the order $J, J-1,\ldots,1,0$, it computes the mapping at a node based on the mapping values on its children in the next resolution. The algorithm achieves the same computational complexity for evaluating the posterior exactly as Mallat's algorithm, that is, linear both in n and in T.

Remark II: The term $\xi_{0,0}(0)$ is the overall marginal likelihood (integrating out all the latent variables) given the hyperparameters, which we can use to set the hyperparameters through

a common empirical Bayes strategy—maximum marginal likelihood estimation (MMLE).

Given the analytical form of the joint posterior, the posterior mean of $\boldsymbol{z}$ can also be computed exactly. To this end, we first use a top-down pyramid algorithm to compute the posterior marginal probability of the hidden states as follows. In the order $j = 1, 2, \ldots, J$, the posterior marginal probability of $S_{j,k}$ for each k is available as

$$\Pr(S_{j,k} = s' \,|\, \mathcal{D}) = \sum_{s \in \{0,1\}} \Pr(S_{j-1,\lfloor k/2 \rfloor} = s \,|\, \mathcal{D}) \cdot \Pr(S_{j,k} = s' \,|\, S_{j-1,\lfloor k/2 \rfloor} = s, \mathcal{D}).$$

Then the posterior mean of $\boldsymbol{z}$ is given by

$$\tilde{z}_{j,k} := \mathrm{E}(z_{j,k} \,|\, \mathcal{D}) = \Pr(S_{j,k} = 1 \,|\, \mathcal{D}) \cdot \frac{n}{n + \tau_j^{-1}} \bar{d}_{j,k}, \tag{7}$$

which has an intuitive explanation in terms of shrinkage. The average of the observed wavelet coefficients $\bar{d}_{j,k}$ is shrunk toward the prior mean 0 with the amount of shrinkage being averaged over the different shrinkage states. By applying an inverse DWT to $\tilde{\boldsymbol{z}}$ we can get the posterior mean of $\boldsymbol{f}$, $E(\boldsymbol{f} \,|\, \mathcal{D}) = W^{-1}\tilde{\boldsymbol{z}}$.

In addition to the posterior mean, one can construct credible intervals for $\boldsymbol{z}$ and for $\boldsymbol{f}$ by sampling from the joint posterior of $\{(S_{j,k}, z_{j,k}, \sigma^2_{j,k}) : j = 0, 2, \ldots, J, k = 0, 1, \ldots, 2^j - 1\}$ according to Theorem 1. Because the exact posterior is available, no MCMC is needed and the sampling is standard Monte Carlo. Given a posterior samples Bayesian inference can proceed as usual. For example, a credible band for $\boldsymbol{f}$ is available from a posterior sample $\boldsymbol{f}$, attained through applying an inverse DWT to the posterior sample on $\boldsymbol{z}$. We illustrate the work of this model in function denoising through simulations in Section 3.1.

2.2 Wavelet fANOVA with normal-inverse-Gamma Markov grove

Next we present our main methodology for fANOVA. We first introduce our framework in one-way fANOVA as the notation is much simpler with all the essential components of the

framework present, and then generalize the formulation to the general multi-way case.

One-way fANOVA. Suppose we have G groups of independent functional observations whose values are attained at T equidistant points $\boldsymbol{y}^{(g,i)} = (y_1^{(g,i)}, \ldots, y_T^{(g,i)})$. Suppose

$$\begin{aligned} \boldsymbol{y}^{(g,i)} &= \boldsymbol{f}^{(g)} + \boldsymbol{\epsilon}^{(g,i)} \\ \boldsymbol{\epsilon}^{(g,i)} &\sim \mathrm{N}(\boldsymbol{0}, \Sigma_\epsilon), \end{aligned} \tag{8}$$

where $g = 1, \ldots, G$ is the group index, $i = 1, \ldots, n_g$ is the index for the replicates in the gth group, and $\Sigma_\epsilon = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_T^2)$. The (one-way) fANOVA problem concerns identifying the variation among $\boldsymbol{f}^{(g)}$, if any, from that in the noise.

Treating the first group as the baseline $\boldsymbol{f} = \boldsymbol{f}^{(1)}$, and letting $\boldsymbol{b}^{(g)} = \boldsymbol{f}^{(g)} - \boldsymbol{f}^{(1)}$ be the contrast of each group to the baseline, we can write

$$\boldsymbol{y}^{(g,i)} = \boldsymbol{f} + \boldsymbol{b}^{(g)} + \boldsymbol{\epsilon}^{(g,i)}. \tag{9}$$

By design $\boldsymbol{b}^{(1)} = \boldsymbol{0}$, and the ANOVA problem boils down to inference on the contrast functions $\boldsymbol{b}^{(g)}$ for $g \geq 2$. After applying the DWT we obtain

$$\boldsymbol{d}^{(g,i)} = \boldsymbol{z} + \boldsymbol{\beta}^{(g)} + \boldsymbol{u}^{(g,i)}$$

where $\boldsymbol{z} = W\boldsymbol{f}$, $\boldsymbol{\beta}^{(g)} = W\boldsymbol{b}^{(g)}$, and $\boldsymbol{u}^{(g,i)} = W\boldsymbol{\epsilon}^{(g,i)}$ for $g = 1, 2, \ldots, G$, and $i = 1, 2, \ldots, n_g$. We can again write the model in a node-specific manner:

$$d_{j,k}^{(g,i)} = z_{j,k} + \beta_{j,k}^{(g)} + u_{j,k}^{(g,i)} \quad \text{where } u_{j,k}^{(g,i)} \overset{\text{ind}}{\sim} \mathrm{N}(0, \sigma_{j,k}^2).$$

As before, we introduce a latent indicator $S_{j,k}$ for each $z_{j,k}$ and adopt the same NIG-MT setup for $\{z_{j,k}, S_{j,k}, \sigma_{j,k}^2 : (j,k) \in \mathcal{T}\}$ as given in Eqs. (2), (3), (4), and (6). Under this formulation, ANOVA can be accomplished through inference on the contrast coefficients $\beta_{j,k}^{(g)}$. To this end, we introduce another latent indicator $R_{j,k} \in \{0, 1\}$ such that

$$[\beta_{j,k}^{(2)}, \ldots, \beta_{j,k}^{(G)} | \sigma_{j,k}^2, R_{j,k}] \overset{\text{iid}}{\sim} \mathrm{N}(0, R_{j,k} \cdot \upsilon_j \sigma_{j,k}^2) \tag{10}$$

where similar to τ_j, $\upsilon_j = 2^{-\alpha j}\upsilon$ is a scaling parameter that characterizes the size of the differences across the factor levels. Our motivation to defining a different scaling parameter υ than τ for the β's is that the β's characterize the difference across the factor levels while the z that for the baseline function mean. It is often the case that the scale of the z's are substantially different than that of the β's.

Just as $\{S_{j,k} : (j,k) \in \mathcal{T}\}$ are modeled in a correlated manner to capture the spatial-scale dependency in $z_{j,k}$, we do that for the $\{R_{j,k} : (j,k) \in \mathcal{T}\}$ as well. Intuitively, if a factor contributes to the variation at one location-scale combination, then it typically contributes to the variation at the children/neighbor nodes as well. Again, an MT is a convenient choice for jointly modeling the latent indicators $\{R_{j,k} : (j,k) \in \mathcal{T}\}$. We let $\boldsymbol{\kappa}_{j,k}$ denote the corresponding state transition matrix (or the initial probability vector when $(j,k) = (0,0)$), which can be specified in the same way as given in (5) for $\boldsymbol{\rho}_{j,k}$. That is for $r, r' \in \{0,1\}$,

$$\Pr(R_{j,k} = r' | R_{j-1,\lfloor k/2 \rfloor} = r) = \kappa_{j,k}(r,r') \text{ for } j > 0 \quad \text{and} \quad \Pr(R_{0,0} = r) = \kappa_{0,0}(r). \tag{11}$$

Now we arrive at a fully specified joint model on $\{z_{j,k}, S_{j,k}, \boldsymbol{\beta}_{j,k}, R_{j,k}, \sigma^2_{j,k} : (j,k) \in \mathcal{T}\}$ given by Eqs. (2), (3), (4), (6), (10), and (11). It is specified by the NIG conjugate priors on $(z_{j,k}, \boldsymbol{\beta}_{j,k}, \sigma^2_{j,k})$ given the latent indicators, and two MTs on the latent indicators. For this reason, we shall refer to this model as a *normal-inverse-Gamma Markov grove* (NIG-MG).

Next we show how Bayesian inference can be carried out for the NIG-MG. It turns out that the joint posterior $\{z_{j,k}, \boldsymbol{\beta}_{j,k}, S_{j,k}, R_{j,k}, \sigma^2_{j,k} : (j,k) \in \mathcal{T}\}$ can again be computed analytically through a pyramid algorithm whose complexity is linear in both n and T. Accordingly, posterior marginal and joint null/alternative probabilities can also be evaluated exactly, and one can sample from the exact posterior using standard Monte Carlo.

We write the node-specific model in matrix notation:

$$\boldsymbol{d}_{j,k} = X\boldsymbol{\theta}_{j,k} + \boldsymbol{u}_{j,k},$$

where $\boldsymbol{d}_{j,k} = (d_{j,k}^{(1,1)}, \ldots, d_{j,k}^{(n_G,G)})'$ is the vector of the wavelet coefficients for all the observations at node (j,k), $\boldsymbol{\theta}_{j,k} = (z_{j,k}, \beta_{j,k}^{(2)}, \ldots, \beta_{j,k}^{(G)})'$ is the vector of the wavelet coefficients for the mean functions, $\boldsymbol{u}_{j,k} = (u_{j,k}^{(1,1)}, \ldots, u_{j,k}^{(n_G,G)})'$ is the vector of the residual errors, and X is the design matrix. The design matrix can be written as

$$X = (\mathbb{1}_n, \boldsymbol{e}_2, \ldots, \boldsymbol{e}_G),$$

where $n = \sum_{g=1}^{G} n_g$, $\mathbb{1}_n$ is a vector of n ones, and $\boldsymbol{e}_g$ is a binary vector where the hth element is equal to one if the hth observation belongs to group g, and equal to zero otherwise. We also define the following matrices for $s, r \in \{0,1\}$:

$$X(s,r) = (s\mathbb{1}_n, r\boldsymbol{e}_2, \ldots, r\boldsymbol{e}_G), \quad \Lambda_j = \mathrm{diag}(1/\tau_j, \underbrace{1/\upsilon_j, 1/\upsilon_j, \ldots, 1/\upsilon_j}_{G-1 \text{ copies}}),$$

and

$$M(s,r) = \begin{pmatrix} s & sr\mathbb{1}'_{G-1} \\ sr\mathbb{1}_{G-1} & r\mathbb{1}_{G-1}\mathbb{1}'_{G-1} \end{pmatrix}.$$

The marginal likelihood for the node-specific model on (j,k) given $S_{j,k} = s$ and $R_{j,k} = r$ is

$$m_{j,k}(s,r) = \frac{(\nu\sigma_0^2)^{\nu+1}\Gamma(\nu + n/2 + 1)}{(2\pi)^{n/2}\Gamma(\nu+1)} \cdot \frac{|\Lambda_j|^{1/2}}{|\Lambda_j^*(s,r)|^{1/2}} \cdot \left[\nu\sigma_0^2 + \Upsilon_{j,k}(s,r)\right]^{-\nu-n/2-1}. \tag{12}$$

where

$$\Upsilon_{j,k}(s,r) = \left\{\boldsymbol{d}'_{j,k}\boldsymbol{d}_{j,k} - [\boldsymbol{\mu}_{j,k}^*(s,r)]'\Lambda_j^*(s,r)\boldsymbol{\mu}_{j,k}^*(s,r)\right\}/2,$$
$$\Lambda_j^*(s,r) = X(s,r)'X(s,r) + \Lambda_j, \quad \text{and} \quad \boldsymbol{\mu}_{j,k}^*(s,r) = [\Lambda_j^*(s,r)]^{-1}[X(s,r)'\boldsymbol{d}_{j,k}].$$

Theorem 2. The joint posterior on $\{z_{j,k}, \boldsymbol{\beta}_{j,k}, S_{j,k}, R_{j,k}, \sigma_{j,k}^2 : (j,k) \in \mathcal{T}\}$ is as follows.

- The marginal posterior of the hidden states $\{(S_{j,k}, R_{j,k}) : (j,k) \in \mathcal{T}\}$ is an MT defined on the product state-space $\{0,1\} \times \{0,1\}$ with

  1. State transition probabilities:

$$\Pr(S_{j,k} = s', R_{j,k} = r' | S_{j-1,\lfloor k/2 \rfloor} = s, R_{j-1,\lfloor k/2 \rfloor} = r, \mathcal{D}) = \rho_{j,k}(s, s')\kappa_{j,k}(r, r')\phi_{j,k}(s', r')/\xi_{j,k}(s, r),$$

for $j = 1, \ldots, J$.

2. Initial state probabilities:

$$\Pr(S_{0,0} = s, R_{0,0} = r | \mathcal{D}) = \rho_{0,0}(s)\kappa_{0,0}(r)\phi_{0,0}(s, r)/\xi_{0,0}(0, 0).$$

- The conditional posterior of $\sigma^2_{j,k}$ given $S_{j,k}$ and $R_{j,k}$ is:

$$[\sigma^2_{j,k} | S_{j,k}, R_{j,k}, \mathcal{D}] \sim \text{Inv-Gamma}\left(\nu + 1 + \frac{n}{2}, \nu\sigma_0^2 + \Upsilon_{j,k}(S_{j,k}, R_{j,k})\right).$$

- The posterior of $z_{j,k}, \beta^{(2)}_{j,k}, \ldots, \beta^{(G)}_{j,k}$ given $S_{j,k}$, $R_{j,k}$ and $\sigma^2_{j,k}$ is given as follows

$$[z_{j,k}, \beta^{(2)}_{j,k}, \ldots, \beta^{(G)}_{j,k} | \sigma^2_{j,k}, S_{j,k}, R_{j,k}, \mathcal{D}] \\ \sim \text{N}\left(\boldsymbol{\mu}^*_{j,k}(S_{j,k}, R_{j,k}), \sigma^2_{j,k}\, M(S_{j,k}, R_{j,k}) \circ [\Lambda^*_j(S_{j,k}, R_{j,k})]^{-1}\right).$$

where $\circ$ represents the Hadamard product, and for any matrix A. The mappings $\phi_{j,k}(s, r)$ and $\xi_{j,k}(s, r) : \{0, 1\} \times \{0, 1\} \to [0, +\infty)$ can be computed recursively through a pyramid algorithm as follows:

$$\phi_{j,k}(s, r) = \begin{cases} m_{j,k}(s, r) \cdot \xi_{j+1,2k}(s, r) \cdot \xi_{j+1,2k+1}(s, r) & \text{for } j = 0, 1, 2 \ldots, J-1 \\ m_{j,k}(s, r) & \text{for } j = J, \end{cases}$$

$$\xi_{j,k}(s, r) = \begin{cases} \sum_{s',r'} \rho_{j,k}(s, s') \cdot \kappa_{j,k}(r, r')\phi_{j,k}(s', r') & \text{for } j = 1, 2, \ldots, J \\ \sum_{s',r'} \rho_{0,0}(s') \cdot \kappa_{0,0}(r') \cdot \phi_{0,0}(s', r') & \text{for } j = 0. \end{cases}$$

Once the joint posterior is computed following the theorem, Bayesian inference can proceed in the usual manner. In particular, for testing the presence of a variance contribution from the factor, $R_{j,k}$ is an indicator for whether the null hypothesis at location-scale (j, k):

$$H_{j,k} : \beta^{(1)}_{j,k} = \beta^{(2)}_{j,k} = \cdots = \beta^{(G)}_{j,k} = 0$$

is false. Thus $P(R_{j,k} = 1 \,|\, \mathcal{D})$ represents the posterior probability for the factor to contribute to the variation at location-scale (j, k). For this reason, we shall refer to $P(R_{j,k} = 1 \,|\, \mathcal{D})$ as the *posterior marginal alternative probability* (PMAP) for location-scale (j, k). Next we show how to compute PMAPs through a top-down pyramid algorithm.

Corollary 1. *For $j = 1, 2, \ldots, J$, the posterior marginal distribution of $(S_{j,k}, R_{j,k})$ can be computed recursively as*

$$\begin{aligned}
&\Pr(S_{j,k} = s', R_{j,k} = r' \,|\, \mathcal{D}) \\
&= \sum_{s,r} \Pr(S_{j-1,\lfloor k/2 \rfloor} = s, R_{j-1,\lfloor k/2 \rfloor} = r \,|\, \mathcal{D}) \times \Pr(S_{j,k} = s', R_{j,k} = r' | S_{j-1,\lfloor k/2 \rfloor} = s, R_{j-1,\lfloor k/2 \rfloor} = r, \mathcal{D})
\end{aligned}$$

using the initial and transition probabilities given in Theorem 2. Then the PMAPs are given by $\Pr(R_{j,k} = 1 \,|\, \mathcal{D}) = \sum_s \Pr(S_{j,k} = s, R_{j,k} = r \,|\, \mathcal{D})$.

In the next corollary, we show how to compute the posterior probability for the presence of factor effects at *any* (i.e., at least one) location-scale combination. This probability, which we refer to as the posterior joint alternative probability (PJAP) can be used for testing the "global" null hypothesis that the factor does not contribute to the variation at all.

Corollary 2. *For all $(j, k) \in \mathcal{T}$ and $s \in \{0, 1\}$, let*

$$\varphi_{j,k}(s) = P(R_{j',k'} = 0 \text{ for all } (j', k') \in \mathcal{T}(j, k) \,|\, S_{j-1,\lfloor k/2 \rfloor} = s, R_{j-1,\lfloor k/2 \rfloor} = 0, \mathcal{D})$$

where $\mathcal{T}(j, k)$ denotes the subtree in $\mathcal{T}$ with (j, k) as the root, i.e., $\mathcal{T}(j, k)$ includes (j, k) and all of its descendants in $\mathcal{T}$. Then we can compute $\varphi_{j,k}(s)$ by the following pyramid algorithm

$$\varphi_{j,k}(s)$$
$$= \begin{cases} \sum_{s'} \Pr(S_{j,k}=s', R_{j,k}=0 | S_{j-1,\lfloor k/2 \rfloor}=s, R_{j-1,\lfloor k/2 \rfloor}=0, \mathcal{D}) \varphi_{j+1,2k}(s') \varphi_{j+1,2k+1}(s') & \text{for } j < J \\ \sum_{s'} \Pr(S_{j,k}=s', R_{j,k}=0 | S_{j-1,\lfloor k/2 \rfloor}=s, R_{j-1,\lfloor k/2 \rfloor}=0, \mathcal{D}) & \text{for } j = J. \end{cases}$$

The posterior joint null probability (PJNP) is given by

$$\Pr(R_{j,k}=0 \text{ for all } (j,k) \in \mathcal{T} \,|\, \mathcal{D}) = \varphi_{0,0}(0).$$

Accordingly, the posterior joint alternative probability (PJAP) is $1-\varphi_{0,0}(0)$.

Remark: Corollary 1 and 2 can also be applied to the prior model to get the prior marginal alternative probabilities, which can be used to elicit the prior specification on the hyperparameters. See Section 2.3 for more details.

Theorem 2 allows us to draw posterior samples of $\{z_{j,k}, \boldsymbol{\beta}_{j,k}, S_{j,k}, R_{j,k}, \sigma^2_{j,k} : (j,k) \in \mathcal{T}\}$ using standard Monte Carlo (not MCMC). Based on this posterior sample, we can also complete other inference tasks such as computing the posterior mean of $\boldsymbol{b}^{(g)}$ and constructing credible bands for $\boldsymbol{b}^{(g)}$ which quantifies the posterior uncertainty of the factor contribution to the functional variation. This can be achieved by applying inverse DWT to the posterior draws of $\boldsymbol{\beta}_{j,k}$'s. We will illustrate this in the numerical examples.

Multi-way fANOVA. The NIG-MG model for one-way fANOVA can be naturally extended to the case with multiple factors by specifying one MT for each factor to capture the location-scale clustering of each factor effect. The complication is mainly in the notation.

Suppose now we have L factors, and the lth factor has G_l levels. Now suppose for each factor combination $\boldsymbol{g} = (g_1, g_2, \ldots, g_L) \in \{1,2,\ldots,G_1\} \times \{1,2,\ldots,G_2\} \times \cdots \times \{1,2,\ldots,G_L\}$, we have $n_{\boldsymbol{g}}$ independent functional observations $\boldsymbol{y}^{(\boldsymbol{g},i)} = (y_1^{\boldsymbol{g},i}, y_2^{\boldsymbol{g},i}, \ldots, y_T^{\boldsymbol{g},i})$ for $i = 1,2,\ldots,n_{\boldsymbol{g}}$,

and suppose each observation arise from the following model

$$\begin{aligned} \boldsymbol{y}^{(\boldsymbol{g},i)} &= \boldsymbol{f}_1^{(g_1)} + \boldsymbol{f}_2^{(g_2)} + \cdots + \boldsymbol{f}_L^{(g_L)} + \boldsymbol{\epsilon}^{(\boldsymbol{g},i)} \\ \boldsymbol{\epsilon}^{(\boldsymbol{g},i)} &\sim \mathrm{N}(\mathbf{0}, \Sigma_\epsilon), \end{aligned} \tag{13}$$

where $g = 1, \ldots, G$ is the group index, $i = 1, \ldots, n_g$ is the index for the replicates in the gth group, and $\Sigma_\epsilon = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_T^2)$.

Now let $\boldsymbol{f} = \sum_{l=1}^{L} \boldsymbol{f}_l^{(1)}$ and $\boldsymbol{b}_l^{(g_l)} = \boldsymbol{f}_l^{(g_l)} - \boldsymbol{f}_l^{(1)}$ for $g_l = 1, 2, \ldots, G_l$. Then

$$\boldsymbol{y}^{(\boldsymbol{g},i)} = \boldsymbol{f} + \sum_{l=1}^{L} \boldsymbol{b}_l^{(g_l)} + \boldsymbol{\epsilon}^{(\boldsymbol{g},i)}.$$

After applying the DWT we obtain:

$$\boldsymbol{d}^{(\boldsymbol{g},i)} = \boldsymbol{z} + \boldsymbol{\beta}_l^{(g_l)} + \boldsymbol{u}^{(\boldsymbol{g},i)}$$

where $\boldsymbol{z} = W\boldsymbol{f}$, $\boldsymbol{\beta}_l^{(g_l)} = W\boldsymbol{b}_l^{(g_l)}$, and $\boldsymbol{u}^{(\boldsymbol{g},i)} = W\boldsymbol{\epsilon}^{(\boldsymbol{g},i)}$ for $\boldsymbol{g}$ and $i = 1, 2, \ldots, n_{\boldsymbol{g}}$. Again the corresponding node-specific model is

$$d_{j,k}^{(\boldsymbol{g},i)} = z_{j,k} + \sum_{l=1}^{L} \beta_{l\,j,k}^{(g_l)} + u_{j,k}^{(\boldsymbol{g},i)} \quad \text{where } u_{j,k}^{(\boldsymbol{g},i)} \overset{\text{ind}}{\sim} \mathrm{N}(0, \sigma_{j,k}^2).$$

Just as in the one-way case, in addition to the latent indicators $S_{j,k}$ for $z_{j,k}$, we introduce an indicator $R_{l\,j,k}$ for each factor effect such that $R_{l\,j,k} = 1$ if and only if

$$H_{l\,j,k} : \beta_{l\,j,k}^{(1)} = \beta_{l\,j,k}^{(2)} = \cdots = \beta_{l\,j,k}^{(G_l)} = 0 \quad \text{is false.}$$

Because each factor contribution will be correlated across location and scale, we adopt an MT model on each of $\{S_{j,k} : (j,k) \in \mathcal{T}\}$ and $\{R_{l\,j,k} : (j,k) \in \mathcal{T}\}$ for $l = 1, 2, \ldots, L$, to induce the proper location-scale dependency. This results in a "grove" of $L + 1$ Markov trees,

$$\{S_{j,k} : (j,k) \in \mathcal{T}\} \sim \mathrm{MT}(\boldsymbol{\rho}_{j,k}) \quad \text{and} \quad \{R_{l\,j,k} : (j,k) \in \mathcal{T}\} \sim \mathrm{MT}(\boldsymbol{\kappa}_{j,k}) \quad \text{for } l = 1, 2, \ldots, L,$$

which are independent given the transition matrices $\boldsymbol{\rho}_{j,k}$ and $\boldsymbol{\kappa}_{j,k}$. We specify $\boldsymbol{\rho}_{j,k}$ and $\boldsymbol{\kappa}_{j,k}$ as in Eq. (5). We allow the hyperparameters (η, γ) to be different for $\boldsymbol{\kappa}_{j,k}$ than for $\boldsymbol{\rho}_{j,k}$. This is necessary because the sparsity (as characterized by η) as well as the spatial-

dependency (as characterized by γ) can be very different for the baseline mean function than for factor effects. For example, the baseline function may be a smooth function, while a factor effect consists of spiky disturbances to the baseline. We will investigate these scenarios in the numerical studies. Specifically, we let (η_ρ, γ_ρ) and $(\eta_\kappa, \gamma_\kappa)$ respectively denote the corresponding hyperparameters for $\boldsymbol{\rho}_{j,k}$ and $\boldsymbol{\kappa}_{j,k}$. We discuss prior specification in Section 2.3.

The hyperparameter sharing that we enforce here by specifying the same prior transition matrix $\boldsymbol{\kappa}_{j,k}$ for the factor effects not only helps attain parsimony but is reasonable from a modeling perspective as well. In most applications there is no *prior* reason to believe any factor contribution to have a different spatial-scale dependency pattern than any other. In situations where one indeed has reasons to believe that the spatial dependency is different among the factor effects, we can specify a different prior transition matrix for the the factors.

Finally, we still adopt the NIG specification on $z_{j,k}$, $\beta^{(g_l)}_{l\,j,k}$, and $\sigma^2_{j,k}$ as in Eqs. (2), (3), and (10). In particular, we adopt a different scaling parameter υ_l for each factor effect

$$[\beta^{(2)}_{l\,j,k}, \ldots, \beta^{(G_l)}_{l\,j,k} | \sigma^2_{j,k}, R_{l\,j,k}] \overset{\text{iid}}{\sim} \mathrm{N}(0, R_{l\,j,k} \cdot 2^{-\alpha j} \upsilon_l \cdot \sigma^2_{j,k})$$

because the effect sizes can often be substantially different among the factors. We now have a fully specified joint model, which is again an NIG-MG consisting of $(L+1)$ Markov trees.

Bayesian inference under this model proceeds in exactly the same fashion as that for one-way fANOVA. Specifically, the marginal likelihood of the node-specific model given $S_{j,k} = s$ and $R_{l\,j,k} = r_{l\,j,k}$ for all l, denoted as $m_{j,k}(s, r_1, r_2, \ldots, r_l)$, also takes the same form as before with the design matrix incorporating all of the factor information. The joint posterior is given by a variant of Theorem 2. (See Supplementary Materials S2 for details.) Similarly, the PMAPs for each factor, $\Pr(R_{l\,j,k} = 1 \,|\, \mathcal{D})$, can be computed analytically through pyramid recursion, as well as the PJAP, $\Pr(R_{l\,j,k} = 1 \text{ for some } (j,k) \in \mathcal{T})$, following Corollary 1 and Corollary 2. Samples can be drawn from the joint posterior through standard Monte Carlo.

2.3 Prior specification

The NIG-MG is specified by the following hyperparameters: $\boldsymbol{\theta} = (\alpha, \tau, \boldsymbol{\upsilon}, \sigma_0, \nu, \eta_\rho, \gamma_\rho, \eta_\kappa, \gamma_\kappa)$, where $\boldsymbol{\upsilon} = (\upsilon_1, \upsilon_2, \ldots, \upsilon_L)$. In determining the choice of these hyparameters, it is important to note their interpretations. In particular, the hyperparameters fall into two categories, respectively called the *scaling* parameters and the *sparsity* parameters. We discuss their specification in turn.

Scaling hyparameters. This category consists of $\boldsymbol{\theta}_{scaling} = (\alpha, \tau, \boldsymbol{\upsilon}, \sigma_0, \nu)$, and they characterize the scale (or size) of either the mean or the factor contribution relative to the size of the errors. Their proper specification depends on the underlying smoothness of the function means and factor effects as well as the signal to noise ratio. We recommend setting these parameters in a data-adaptive manner through empirical Bayes (described later).

Sparsity hyparameters. This category consists of the parameters that determine the prior distributions of the latent indicator variable $S_{j,k}$ and $R_{l\,j,k}$. That is, the parameters that specify the state transition matrices of the MTs: $\boldsymbol{\theta}_{sparsity} = (\eta_\rho, \gamma_\rho, \eta_\kappa, \gamma_\kappa)$. In particular, they determine the *a priori* probability for the presence of "signal" at each location-scale combination, and hence they tune the sparsity of the underlying signal, in terms of the proportion of locations-scale combinations with signals. In particular, they determine the quantities such as the prior probability for the null, that there is no factor contribution at all, and the prior expected number of location-scale combinations on which there is a factor contribution, etc. Depending on the inferential goal at hand, different strategies for specifying the sparsity hyparameters can be adopted.

Specifically, if one's goal is for estimation and prediction, such as in signal denoising and image reconstruction, then a simple strategy is again to choose them by some data-fitting criteria such as empirical Bayes. Note that once a data-adaptive approach is taken in choosing the hyparameters, their face-values lose meanings. For example, one cannot interpret the "prior" null probabilities as before because it is chosen based on the data.

In estimation and prediction one does not care so much about the prior interpretation of these parameters as finding the parameters that render the best predictive performance. Thus in such cases we can treat these sparsity hyparameters in the same way as the scaling parameters and choose them through empirical Bayes. If one's goal is hypothesis testing regarding factor effects, however, in order to maintain the "validity" of the test (such as its level from a frequentist perspective or the prior null probability from a Bayesian one), one should not use data to choose the sparsity parameters but should elicit them based on certain prior criteria. For example, one can choose $(\eta_\kappa, \gamma_\kappa)$ such that the prior joint alternative probability is say 50%, as computed by applying Corollary 2 to the prior model.

In practical problems, one can use a hybrid strategy to specify the sparsity parameters. For example, if one is interested in testing the contribution of one or more factors, but not in the baseline, then one can specify $(\eta_\kappa, \gamma_\kappa)$ through prior elicitation, and use empirical Bayes to choose (η_ρ, γ_ρ) which characterizes the baseline structure.

Empirical Bayes by maximum marginal likelihood. We have referred multiple times to empirical Bayes as a strategy for choosing hyparameters. Specifically, a useful by-product of applying Theorems 1 and 2 is the overall marginal likelihood $P(\mathcal{D} \,|\, \boldsymbol{\theta}) = \xi_{0,0}(0)$, which can be computed through the pyramid algorithm. We can thus find the maximum marginal likelihood estimators (MMLE) for the hyperparameters $\hat{\boldsymbol{\theta}} = \text{argmax}_{\boldsymbol{\theta}} P(\mathcal{D} \,|\, \boldsymbol{\theta})$. This optimization can be carried out using standard numerical methods. In our numerical studies, it is completed using the Nelder-Mead algorithm implemented in the `R` function `optim`.

Alternatively we can use prior elicitation to choose $\boldsymbol{\theta}_{sparsity}$ and MMLE to set the scaling parameters $\hat{\boldsymbol{\theta}}_{scaling} = \text{argmax}_{\boldsymbol{\theta}_{scaling}} P(\mathcal{D} \,|\, \boldsymbol{\theta}_{sparsity}, \boldsymbol{\theta}_{scaling})$. We note that strictly speaking, the empirical Bayes strategy makes the uncertainty quantification such as credible intervals "overly confident". But because in the current context the total number of observations nT is typically much much larger than the number of hyperparameters, this impact on the uncertainty quantification is often small.

2.4 Decision rules and multiplicity control

Next we construct decision rules for calling "significant" factor effects. It is natural to reject the joint null hypothesis of no factor effects whatsoever when the PJAP is large for each factor of interest. But how large is large enough? The threshold for PJAP can either be chosen at a specific level such as 80% provided that the prior, especially in terms of the sparsity parameters, is properly calibrated as described in the previous subsection. Alternatively, the PJAP threshold can be determined empirically by resampling strategies such as permutation.

In most applications, it is not only interesting to know whether a factor contributes to the variation at all but to understand the nature of such contribution. For example, what parts of the sample space is affected by that factor and at what scales. To this end, it is useful to consider decision rules for rejecting the node-specific hypotheses $\{H_{j,k} : (j,k) \in \mathcal{T}\}$ directly, and identify those locations-scale combinations (j,k) at which the factor contributes to the variation. With this perspective, a natural decision rule for rejecting $H_{j,k}$ is when the PMAP $P(R_{j,k} = 1 \,|\, \mathcal{D}) > \delta$ for some threshold $\delta \in (0,1)$. Given any threshold δ, the posterior expected *number of false positives* (NFP) regarding the lth factor, i.e., location-scale combinations that are called to be significant in the lth factor but for which the lth factor has no effect is given by $\mathrm{NFP}(\delta) = \sum_{(j,k):P(R_{j,k}=1\,|\,\mathcal{D})>\delta} P(R_{j,k} = 0 \,|\, \mathcal{D})$. Accordingly, the (Bayesian) false discovery rate (FDR) as defined in [16] is given by

$$\mathrm{FDR}(\delta) = \frac{\sum_{(j,k):P(R_{j,k}=1\,|\,\mathcal{D})>\delta} P(R_{j,k} = 0 \,|\, \mathcal{D})}{|\{(j,k) : P(R_{j,k} = 1 \,|\, \mathcal{D}) > \delta\}|}$$

and thus one can choose δ to achieved the desired FDR.

3 Numerical examples

In this section we provide three numerical examples. In the first example we apply the NIG-MT to denoising a single functional observation using the classical scenarios given in [11].

In the second example we illustrate the work of NIG-MG for one-way fANOVA through simulation and compare it to existing wavelet-based fANOVA methods. Finally, we carry out two-way NIG-MG-based fANOVA for a well-known time-series data set, the "orthosis" data. In all the examples we use the Daubechies least-asymmetric orthonormal compactly supported wavelet with 10 vanishing moments.

3.1 Function denoising

In this example we generate synthetic data from the four test functions proposed by [11], namely *blocks*, *bumps*, *doppler* and *heavisine*. In Figure 1 we plot the four functions and the associated mother wavelet coefficients. For each of the four test functions we consider four levels—1, 3, 5, and 7—of the root signal to noise ratio (RSNR): RSNR $= \sqrt{\sum_{t=1}^{T}(f_t - \bar{f})^2/(T-1)}/\sigma$, where $\bar{f} = \sum_t f_t/T$ and $\Sigma_\epsilon = \text{diag}(\sigma^2, \ldots, \sigma^2)$. The observations are taken at $T = 1024$ equidistant points, and for each function and each RSNR level we generate 100 datasets.

In addition to our NIG-MT model, we apply two additional methods, one Bayesian and the other empirical Bayesian, for wavelet shrinkage—namely [3] and [12]—to the simulated data as comparison. [3] is one of the early well-known Bayesian wavelet regression methods, while the two-group emprical Bayes method "EBayesThresh" introduced in [12] is often regarded as the state-of-the-art in Bayesian wavelet denoising. We do not carry out a comprehensive study to the numerous available wavelet shrinkage methods, but include just these two related methods, especially since our main endeavor is for the fANOVA problem. For interested readers, a more extensive comparative study is available in [12], which showed that the `EBayesThresh` approach compares favorably against existing methods.

For each simulated functional observation we apply the three methods. We evaluate the performance of each method using the average (over location) mean square error (AMSE). For NIG-MT, we use empirical Bayes to set all of the hyperparameters. For the method of

[3], we set the two hyperparameters $\alpha = 0.5$ and $\beta = 1$ as recommended in that paper. The method is available in the R package `wavethresh`. For the EBayesThresh method we use the Laplace prior for the wavelet coefficient under the alternative with the scale parameter of the Laplace set at 0.5, which is the default value given in the `R` package `EBayesThresh`. Both `wavethresh` and `EBayesThresh` are available on `CRAN`.

Table 1 presents AMSE for each of the methods for all four signal functions at the four RSNR levels. The NIG-MT method outperforms the other methods in all simulation scenarios. The performance gain for doppler, bumps, and blocks are more substantial than that for heavisine, which is as expected because the wavelet coefficients for the first three functions display stronger spatial-scale dependency, and so exploiting such dependency is most rewarding.

3.2 Identifying functional variation

Next we carry out a simulation study to evaluate the performance of the NIG-MG method in fANOVA problems. In particular, we simulate from the one-way fANOVA model given in Eq. (8) and (9). We consider the case with $G = 3$ groups and each group has $n_g = 3$ replicate observations. Following the notation in Eq. (9), we let $\boldsymbol{f}$ be the baseline mean for Group 1, and $\boldsymbol{b}^{(i)}$ be the contrast between Group i and Group 1 for $i = 2, 3$. For simplicity, we let $\boldsymbol{b}^{(2)} = -\boldsymbol{b}^{(3)} = \boldsymbol{b}$. That is, the difference between Group 2 and Group 1 is of the same magnitude as that of Group 3 and Group 1, but of the opposite sign.

As in function denoising, the performance of a fANOVA method depends on the nature of the underlying function mean and factor effect. As such, we consider different scenarios in which $\boldsymbol{f}$ and $\boldsymbol{b}^{(i)}$ are of a variety of natures. The scenarios we consider fall into two broad categories. The first corresponds to the case when the cross-group difference is of a global nature, in the sense that they involve a large number of locations, whereas the second category corresponds to cases where the cross-group difference is local, involving only a very

Doppler	RSNR			
	1	3	5	7
NIG-MT	**5.2(0.7)**	**1.1(0.2)**	**0.47(0.06)**	**0.27(0.03)**
EBayesThresh	6.2(0.9)	1.2(0.2)	0.52(0.08)	0.32(0.03)
Bayes	6.1(0.9)	1.5(0.3)	0.67(0.1)	0.41(0.04)
Heavisine	RSNR			
	1	3	5	7
NIG-MT	**1.8(0.5)**	**0.51(0.1)**	**0.26(0.04)**	**0.17(0.03)**
EBayesThresh	2.8(0.9)	0.64(0.1)	0.29(0.06)	0.18(0.03)
Bayes	2.3(0.8)	0.71(0.1)	0.4(0.06)	0.26(0.04)
Bumps	RSNR			
	1	3	5	7
NIG-MT	**17(2)**	**2.9(0.3)**	**1.3(0.1)**	**0.77(0.06)**
EBayesThresh	24(3)	3.6(0.3)	1.5(0.1)	0.83(0.06)
Bayes	30(3)	5.5(0.5)	2.9(0.2)	2.2(0.2)
Blocks	RSNR			
	1	3	5	7
NIG-MT	**10(1)**	**2.1(0.2)**	**0.94(0.08)**	**0.55(0.05)**
EBayesThresh	12(1)	3(0.2)	1.3(0.1)	0.73(0.07)
Bayes	16(2)	4.3(0.4)	2.2(0.2)	1.4(0.1)

Table 1: AMSE ($\times 10^2$) of four methods—NIG-MT, EBayesThresh [12] and Bayes [3] for four different signal functions at four RSNR levels, estimated from 100 simulations. Standard deviations of the AMSEs are given in parentheses. The lowest AMSE in each setting is in bold font.

small subset of locations. We describe the different simulation scenarios in turn.

Global factor effects. In this case, we allow the baseline mean $\boldsymbol{f}$ and the factor difference $\boldsymbol{b}$ to be any of the four signature functions–doppler,heavisine,bumps,blocks, resulting in a total of 16 possible combinations. For each of the 16 combinations, we simulate 500 data sets, for RSNR=$0.5, 1, 1.5, 2$. For each simulated data set, we also simulate a null data set resulting from setting $\boldsymbol{b} = \boldsymbol{0}$. We apply four methods—NIG-MG, the wfANOVA test [14], the tANOVA test [14], and the wavelet minimax test [2]—to each of the simulated data set (both with a difference and the null set), and construct the ROC curve for the corresponding test statistic of each of the four methods. The wfANOVA is in essence the F-test applied to each location-scale combination. The tANOVA is the location-by-location F-test in the original space. The wavelet minimax test statistic is in essence the sum of squares of the wavelet coefficients after proper thresholding at the fine resolutions. For NIG-MG, we use the PJAP as the test statistic for the existence of a cross-group difference.

Figure 2 and Figure 3 show the matrix of ROC curves for RSNR= 1 and RSNR= 2, respectively. (The ROC curves for RSNR= 0.5 and RSNR= 1.5 look similar and so we defer them to Supplementary Materials S3. See Figure S1 and Figure S2.) The NIG-MG outperforms the other methods by comfortable margins at all four RSNR levels.

Local factor effects. We then consider simulation scenarios in which the factor effects are of a local nature, involving only a small fraction of locations. We still let the baseline mean $\boldsymbol{f}$ to be any of the four signature functions. We let $\boldsymbol{b}$ be 0 for most locations, but for just a small interval, we let it be a constant proportion of $\boldsymbol{f}$. The mean functions for each of the three groups are plotted in the first column of Figure 4.

We again apply the four methods to test for the existence of a cross-group difference. The second to fourth columns of Figure 4 show their ROC curves under three RSNR levels 2, 3, and 4. Again the performance advantage of NIG-MG is even more substantial than in the global difference scenarios. This shows the importance of incorporating the spatial-scale

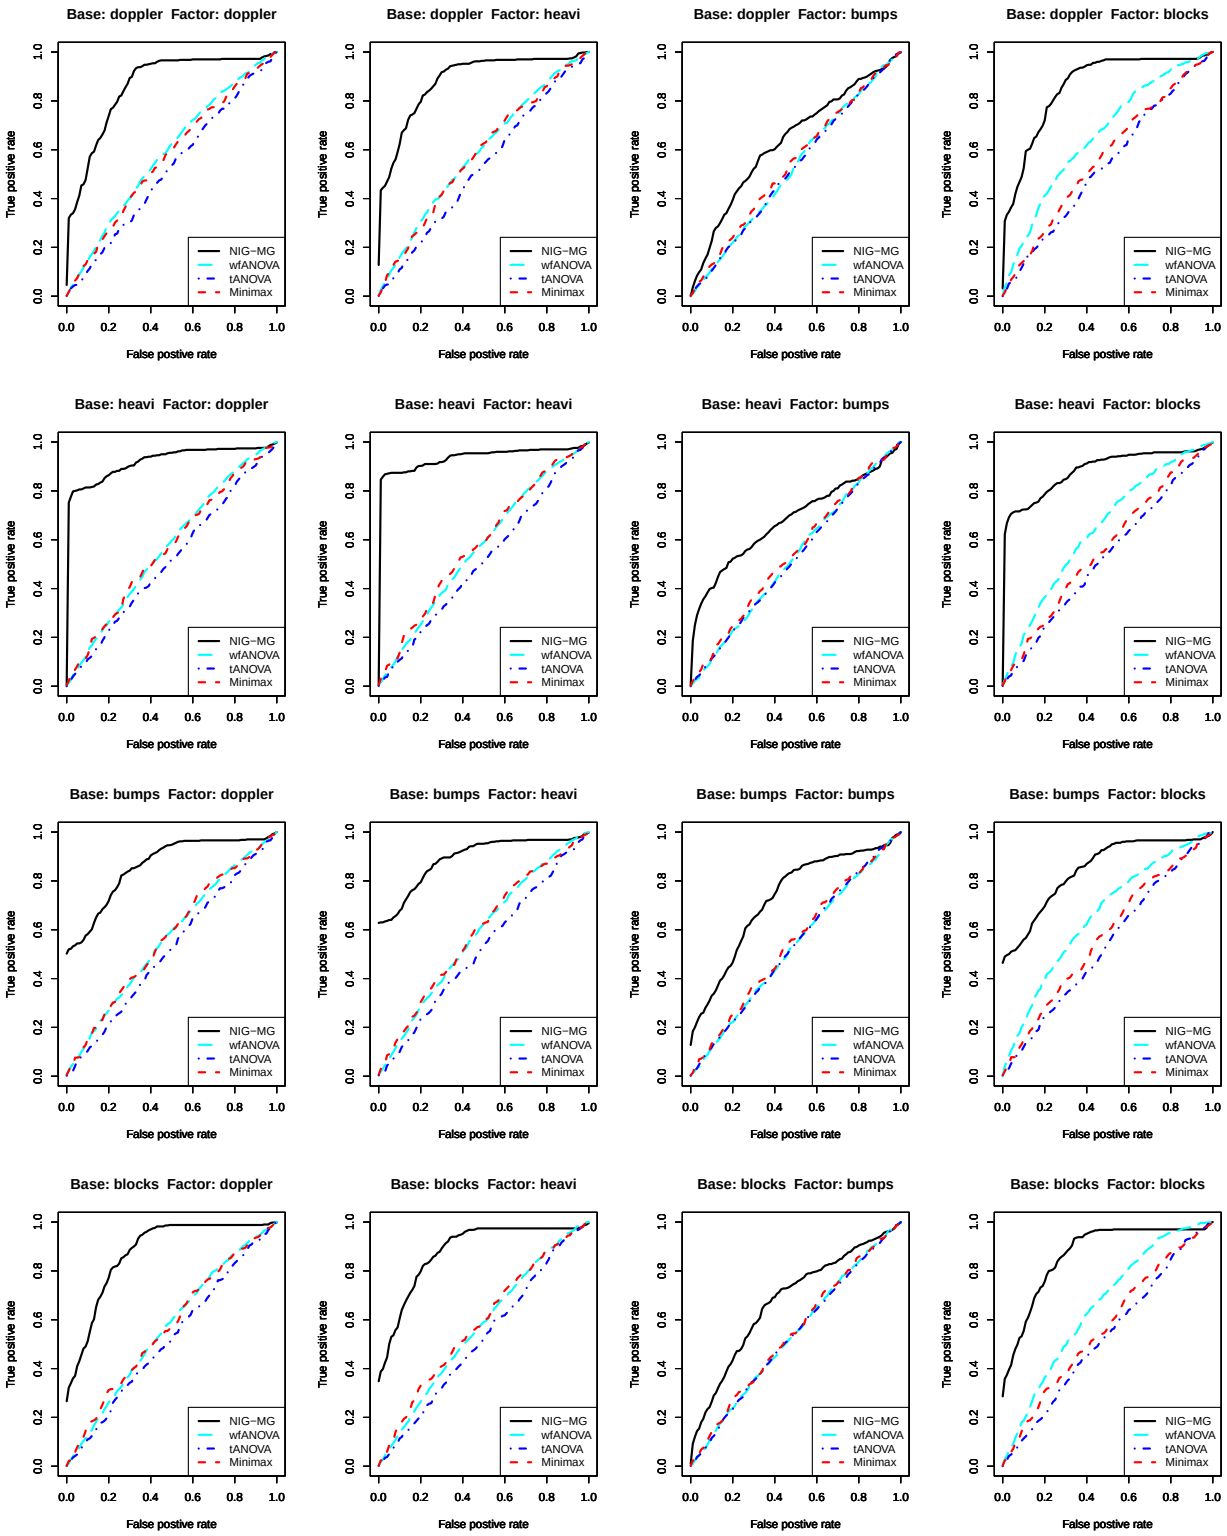

Figure 2: ROC curves (RSNR=1) for four test statistics when the baseline mean and the factor difference are all 16 combinations of the four signature functions.

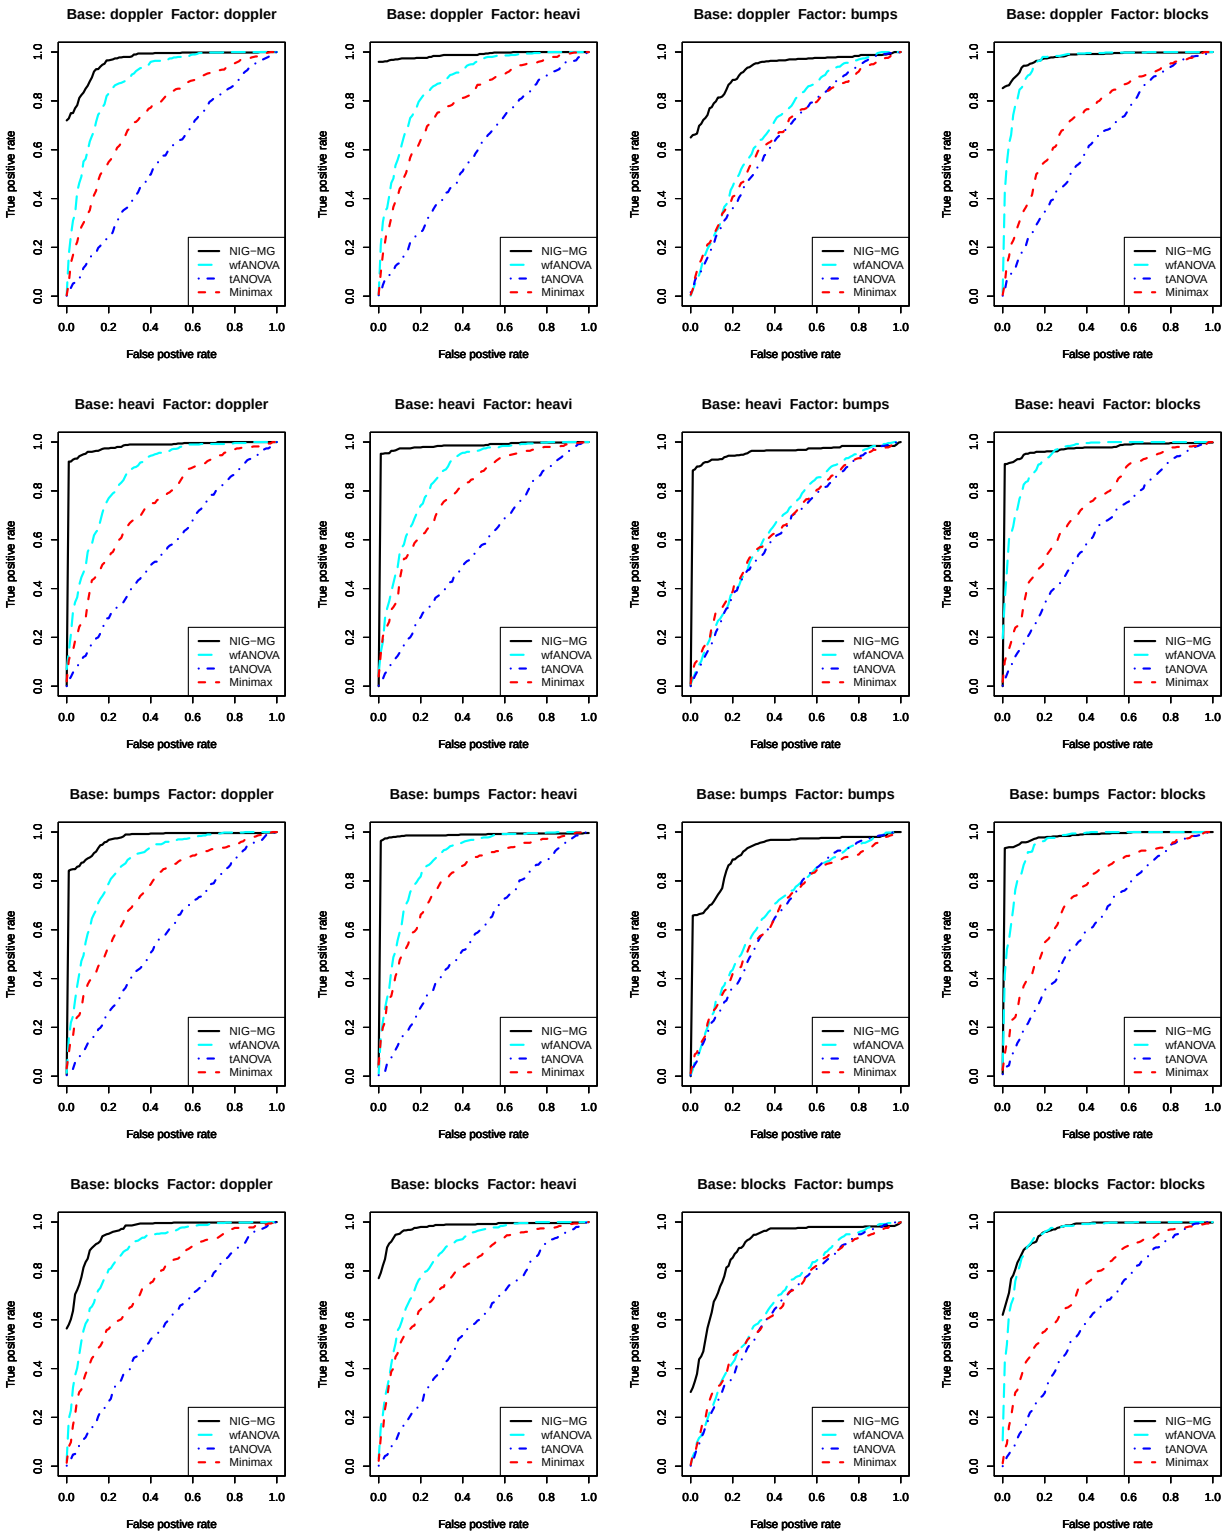

Figure 3: ROC curves (RSNR= 2) for four test statistics when the baseline mean and the factor difference are all 16 combinations of the four signature functions.

dependency when the underlying factor effects is local and so borrowing strength becomes critical to effectively identify such structures.

3.3 Orthosis Dataset

We apply our NIG-MG model to analyze the *orthosis dataset*, a publicly available data original collected by Dr. David Amarantini and Dr. Luc Martin from the Laboratoire Sport et Performance Motrice, Grenoble University, France. This data has been used by several authors as a test-bed for functional data analysis methods [2, 1, 4, 20]. The purpose of the study was to understand the effect of different types of constraints to the knee on movement generation. In the study 7 individuals (i.e., the subjects) wore spring-loaded orthosis on the right knee while stepping in place. Four experimental conditions were considered: a control condition (without orthosis), an orthosis condition (with the orthosis only), and two different springs loaded to the orthosis (spring 1 and spring 2). Ten replicated data sets were collected for each subject under each of the four conditions. The resultant moment for each trial was computed at $T = 256$ equidistant time points. Figure 5 presents the entire data set. We refer the reader to [6] for further detail on the experiment.

A key question to address from this data is how the four experimental conditions result in different knee movement as measured by the functional shape of the measured moments. While the subject-to-subject variation is not of direct interest, it is substantial and must be properly taken into account. Treating both the experimental conditions and the subjects as factors, the experiment corresponds to a two-way ANOVA design. We apply our NIG-MG model for two-way fANOVA. We set the prior null probability to about 50% with $\eta = 0.3$ and $\gamma = 0.4$. The posterior probability for the joint null that there is no difference among the four conditions is virtually zero. Figure 6 presents the PMAPs for the two factors.

While there is strong evidence for difference across the experiment conditions (as well as the subjects), one may also be interested in investigating the factor contributions with regard

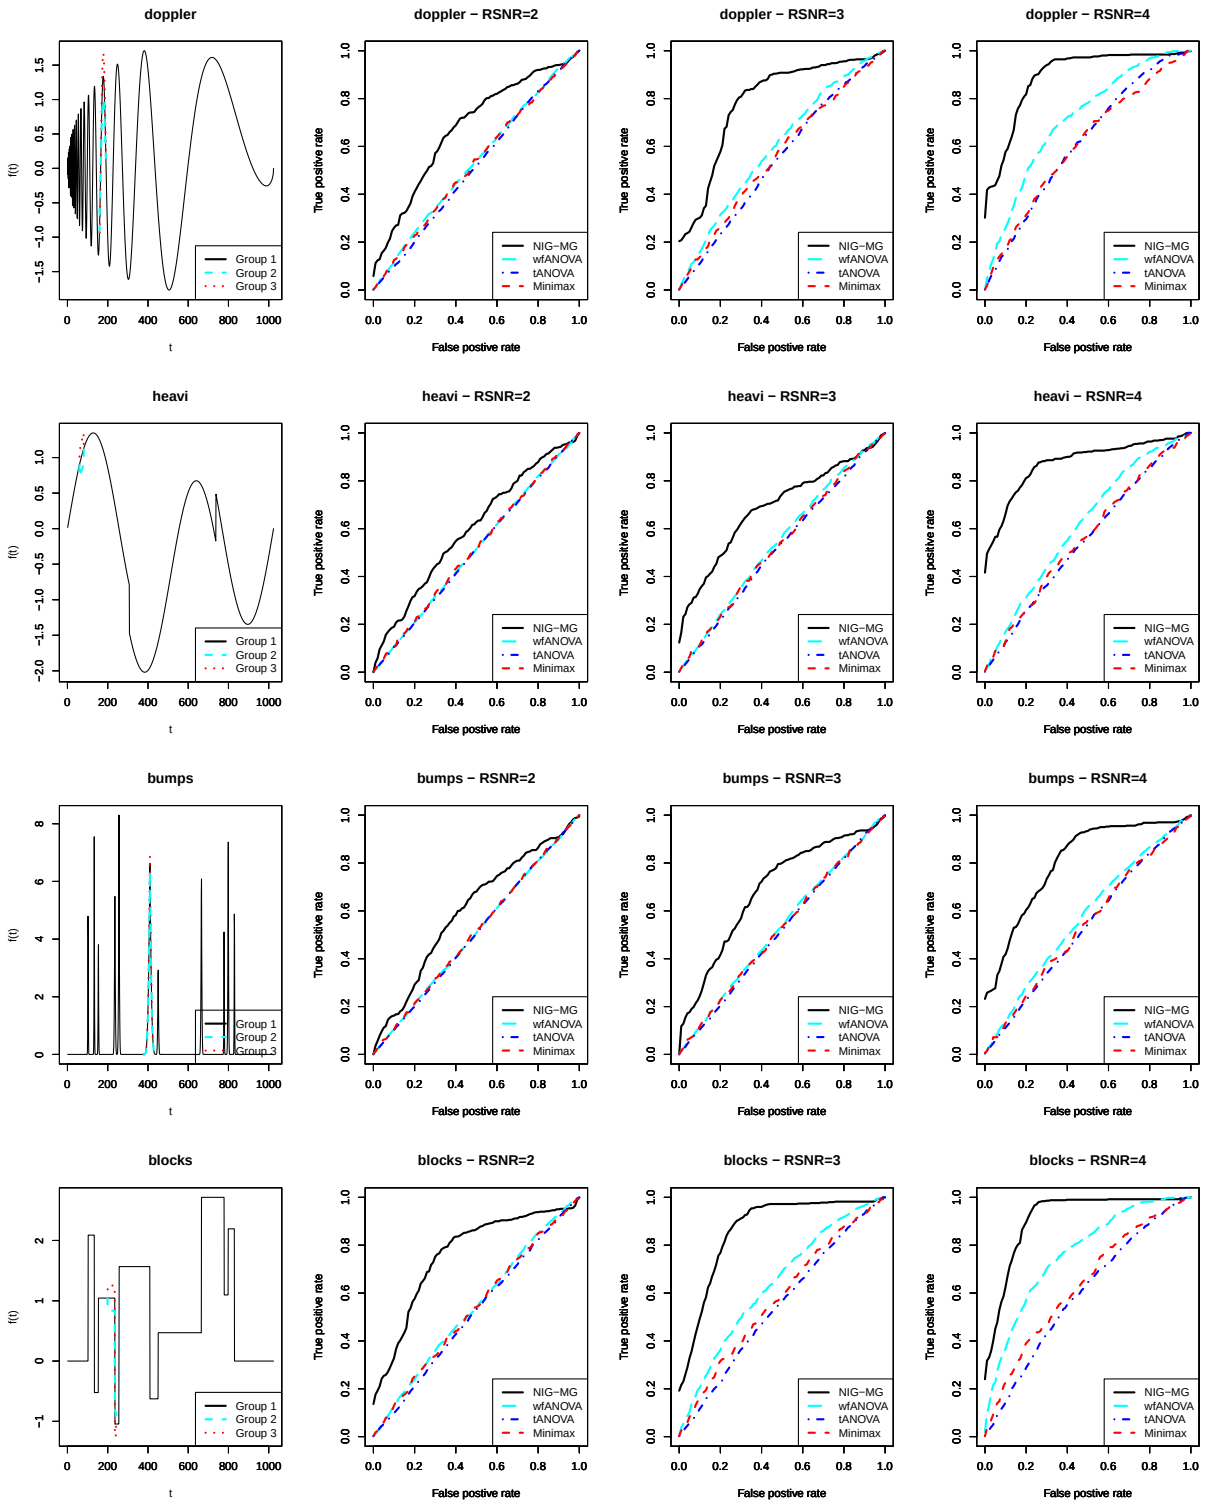

Figure 4: ROC curves for four test statistics under the local cross-group difference scenarios. The rows correspond to different baseline mean functions—first: doppler; second: heavisine; third: bumps; four: blocks. The first columns shows the true mean functions of the three groups. The second to fourth columns correspond to RSNR=$2, 3, 4$ respectively.

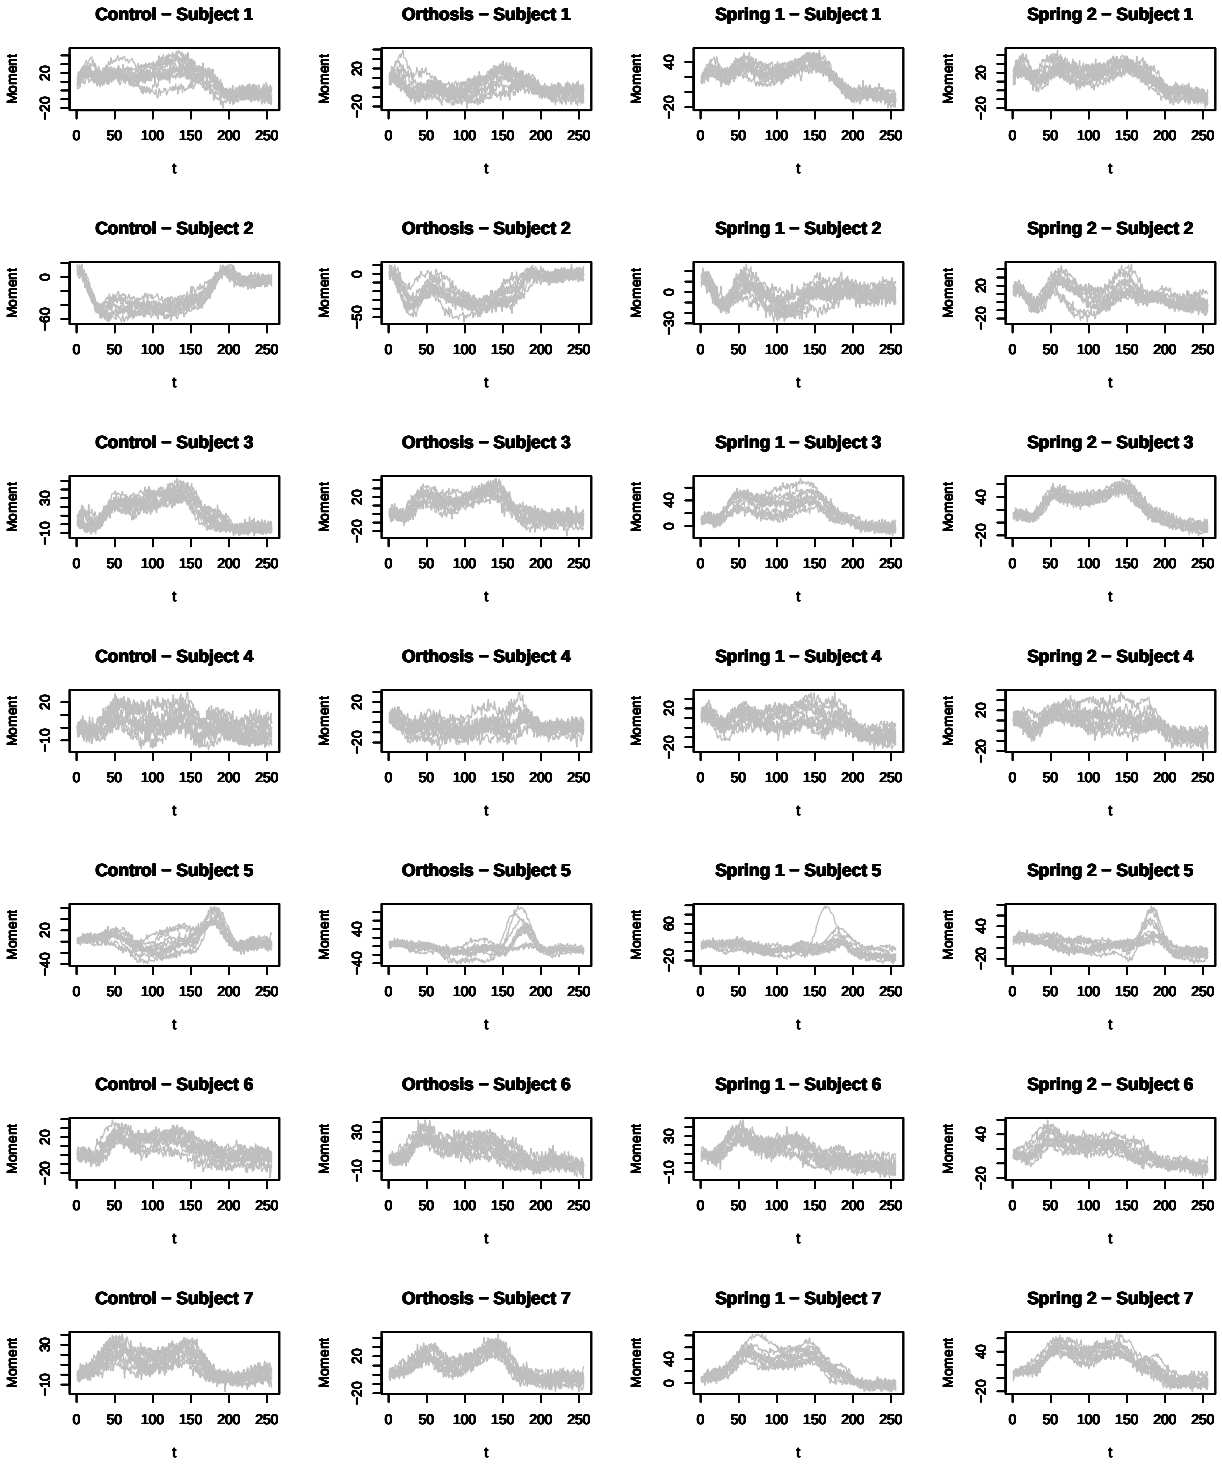
Control – Subject 1
Orthosis – Subject 1
Spring 1 – Subject 1
Spring 2 – Subject 1
Control – Subject 2
Orthosis – Subject 2
Spring 1 – Subject 2
Spring 2 – Subject 2
Control – Subject 3
Orthosis – Subject 3
Spring 1 – Subject 3
Spring 2 – Subject 3
Control – Subject 4
Orthosis – Subject 4
Spring 1 – Subject 4
Spring 2 – Subject 4
Control – Subject 5
Orthosis – Subject 5
Spring 1 – Subject 5
Spring 2 – Subject 5
Control – Subject 6
Orthosis – Subject 6
Spring 1 – Subject 6
Spring 2 – Subject 6
Control – Subject 7
Orthosis – Subject 7
Spring 1 – Subject 7
Spring 2 – Subject 7
Moment
t

Figure 5: The orthosis data. Each row corresponds to a subject, each column to an experimental condition, and each curve is a replicate.

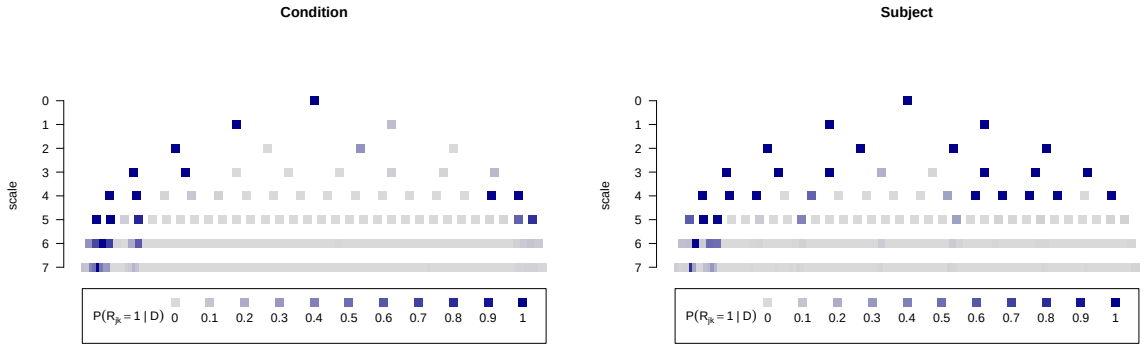

Figure 6: The PMAPs for both factors—the experiment condition (left) and the subject (right) under the NIG-MG model.

to certain contrasts—e.g., the difference between the orthosis only condition and the control, the difference between the two different springs, and the difference between the spring vs no spring conditions. The fully probabilistic nature of the NIG-MG framework allow us to address such tasks while properly taking into account the uncertainty involved through the standard Bayesian recipe—sampling from the joint posterior and construct credible bands for the corresponding contrasts.

Figure 7 presents the credible bands for these three contrasts. In particular, for each contrast we construct two credible bands. The first credible band (shown in the left column in the figure) is constructed using the posterior samples for just the mother wavelet coefficients, excluding the scale coefficient. The quantify the uncertainty in the shape of the functional contrast but not in the mean level of the contrast across all locations (here the time points). In addition, we create also the credible band that uses the posterior samples for both the mother and the scaling coefficients, which incorporates both the functional shape and the mean while taking into account the uncertainty from both sources. According to our knowledge, all of the previous analysis of the data set using wavelet-based methods [2, 1, 4] only provides point estimate of the contrasts without providing uncertainty quantification.

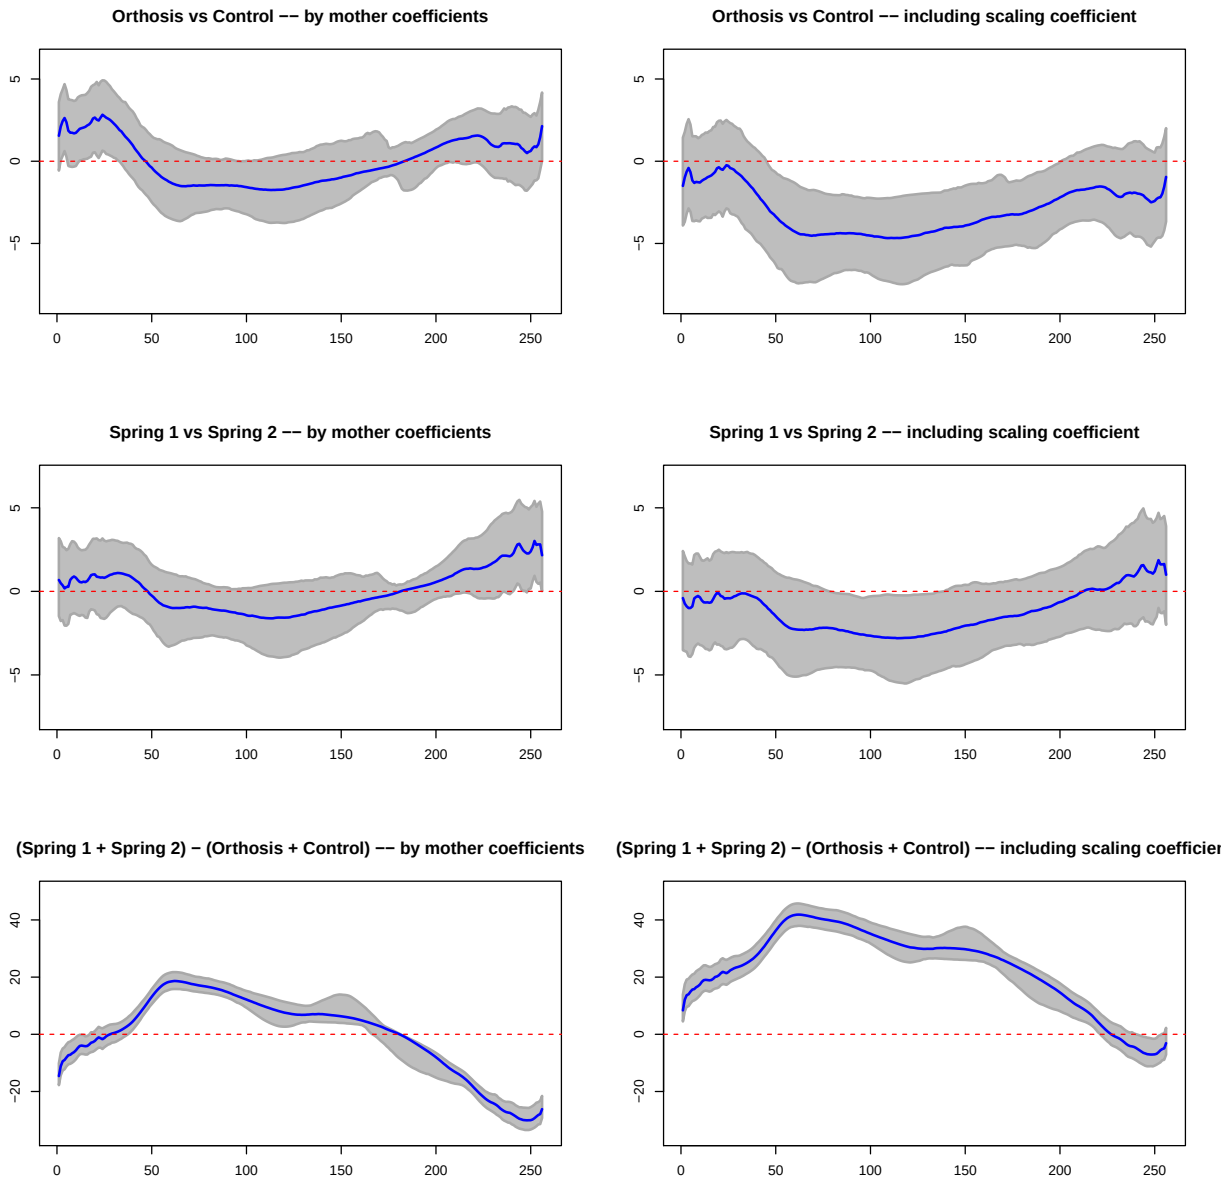

Figure 7: Pointwise credible bands for three different contrasts—Row 1: Orthosis versus Control, Row 2: Spring 1 vs Spring 2, Row 3: Spring vs No spring. Left column: Credible bands constructed based on the posterior samples of the mother wavelet coefficients from the NIG-MG model. Right column: Credible bands constructed based on the posterior samples of the mother wavelet coefficients fro the NIG-MG model as well as posterior draws for the scaling (father wavelet) coefficient.

4 Conclusion

We have introduced a new Bayesian hierarchical model in the wavelet domain for addressing the functional analysis of variance problem. By incorporating a graphical model that links the presence and absence of factor effects on the wavelet coefficients, this model allos effective borrowing of information across locations and scales, and this in turn leads substantial performance gain over methods that ignore such dependency, especially in situations where the underlying factor effects are local. Moreover, the exact posterior of the model can be computed exactly through an efficient pyramid type recursive algorithm which is linear in both the number of observations and the number of locations, i.e., the complexity as Mallat's pyramid algorithm for DWT. In addition, the fully probabilistic nature of the model allows inference to be carried out in a principled manner—uncertainty quantification is achieved through posterior probabilities and credible bands.

The computational complexity of the pyramid algorithm for evaluating the exact posterior of NIG-MG is linear in both n and T. So inference scales well with the number of functional observations n as well as the number of locations m. However, as the number of factors L grows, the computation scales as $O(4^L)$ and so can become infeasible if the number of factors L is large. In most applications of fANOVA, however, the number of factors L is usually small ≤ 5 for which exact inference can be completed quickly.

References

[1] Abramovich, F. and Angelini, C. (2006). Testing in mixed-effects fanova models. *Journal of statistical planning and inference* **136**, 12, 4326–4348.

[2] Abramovich, F., Antoniadis, A., Sapatinas, T., and Vidakovic, B. (2004). Optimal testing in a fixed-effects functional analysis of variance model. *International Journal of Wavelets, Multiresolution and Information Processing* **2**, 04, 323–349.

[3] Abramovich, F., Sapatinas, T., and Silverman, B. W. (1998). Wavelet thresholding via a bayesian approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **60**, 4, 725–749.

[4] Antoniadis, A. and Sapatinas, T. (2007). Estimation and inference in functional mixed-effects models. *Computational Statistics & Data Analysis* **51**, 10, 4793–4813.

[5] Brown, P. J., Fearn, T., and Vannucci, M. (2001). Bayesian wavelet regression on curves with application to a spectroscopic calibration problem. *Journal of the American Statistical Association* **96**, 454, 398–408.

[6] Cahouët, V., Luc, M., and David, A. (2002). Static optimal estimation of joint accelerations for inverse dynamics problem solution. *Journal of Biomechanics* **35**, 11, 1507–1513.

[7] Chipman, H. A., Kolaczyk, E. D., and McCulloch, R. E. (1997). Adaptive bayesian wavelet shrinkage. *Journal of the American Statistical Association* **92**, 440, 1413–1421.

[8] Clyde, M. and George, E. I. (2000). Flexible empirical bayes estimation for wavelets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **62**, 4, 681–698.

[9] Clyde, M., Parmigiani, G., and Vidakovic, B. (1998). Multiple shrinkage and subset selection in wavelets. *Biometrika* **85**, 2, 391–401.

[10] Crouse, M. S., Nowak, R. D., and Baraniuk, R. G. (1998). Wavelet-based statistical signal processing using hidden markov models. *Signal Processing, IEEE Transactions on* **46**, 4, 886–902.

[11] Donoho, D. L. and Johnstone, J. M. (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika* **81**, 3, 425–455.

[12] Johnstone, I. M. and Silverman, B. W. (2004). Needles and straw in haystacks: Empirical bayes estimates of possibly sparse sequences. *Ann. Stat.* **32**, 4 (08), 1594–1649.

[13] Mallat, S. G. (1989). A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.* **11**, 7 (July), 674–693. http://dx.doi.org/10.1109/34.192463.

[14] McKay, J. L., Welch, T. D. J., Vidakovic, B., and Ting, L. H. (2013). Statistically significant contrasts between emg waveforms revealed using wavelet-based functional anova. *Journal of Neurophysiology* **109**, 2, 591–602.

[15] Morris, J. S. and Carroll, R. J. (2006). Wavelet-based functional mixed models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **68**, 2, 179–199.

[16] Müller, P., Parmigiani, G., and Rice, K. (2007). Fdr and bayesian multiple comparisons rules. In *Bayesian Statistics 8*, J. M. Bernardo, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith, and M. West, Eds. Oxford University Press.

[17] Percival, D. and Walden, A. (2006). *Wavelet Methods for Time Series Analysis*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.

[18] Ramsay, J. O. and Silverman, B. W. (2005). *Functional Data Analysis*, 2nd ed. Springer Series in Statistics. Springer.

[19] Rosner, G. L. and Vidakovic, B. (2000). Wavelet functional anova, bayesian false discovery rate, and longitudinal measurements of oxygen pressure in rats.

[20] Zhang, J. (2013). *Analysis of Variance for Functional Data*. Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis Group.

Supplementary Materials

S1. Proofs

Proof of Theorem 1. This theorem is a special case of Theorem 2 and thus follows from Theorem 2. □

Proof of Theorem 2. First define $\mathcal{D}_{j,k}$ to be the data, i.e., the empirical wavelet coefficients observed on all subtree rooted at node (j,k), $\mathcal{T}(j,k)$. Then note that $\xi(s,r)$ is the marginal likelihood on the subtree $\mathcal{T}(j,k)$ given the event that $S_{j-1,\lfloor k/2,\rfloor} = s$ and $R_{j-1,\lfloor k/2,\rfloor} = r$, and $\phi(s,r)$ is the marginal likelihood on $\mathcal{T}(j,k)$ given the event that $S_{j,k} = s$ and $R_{j,k} = r$.

Now by Bayes theorem,

$$\begin{aligned}
&\Pr(S_{j,k} = s', R_{j,k} = r' | S_{j-1,\lfloor k/2\rfloor} = s, R_{j-1,\lfloor k/2\rfloor} = r, \mathcal{D}) \\
=&\frac{\Pr(S_{j,k} = s', R_{j,k} = r', \mathcal{D}_{j,k} | S_{j-1,\lfloor k/2\rfloor} = s, R_{j-1,\lfloor k/2\rfloor} = r)}{\Pr(\mathcal{D}_{j,k} | S_{j-1,\lfloor k/2\rfloor} = s, R_{j-1,\lfloor k/2\rfloor} = r)} \\
=&\frac{\Pr(S_{j,k} = s', R_{j,k} = r' | S_{j-1,\lfloor k/2\rfloor} = s, R_{j-1,\lfloor k/2\rfloor} = r) \cdot \Pr(\mathcal{D}_{j,k} | S_{j,k} = s', R_{j,k} = r')}{\Pr(\mathcal{D}_{j,k} | S_{j-1,\lfloor k/2\rfloor} = s, R_{j-1,\lfloor k/2\rfloor} = r)} \\
=&\rho_{j,k}(s,s')\kappa_{j,k}(r,r')\phi_{j,k}(s',r')/\xi_{j,k}(s,r)
\end{aligned}$$

for $j \geq 1$. Now for $j = 0$, similarly,

$$\begin{aligned}
&\Pr(S_{j,k} = s', R_{j,k} = r' | \mathcal{D}) \\
=&\Pr(S_{j,k} = s', R_{j,k} = r', \mathcal{D}_{j,k}) / \Pr(\mathcal{D}_{j,k}) \\
=&\Pr(S_{j,k} = s', R_{j,k} = r') \cdot \Pr(\mathcal{D}_{j,k} | S_{j,k} = s', R_{j,k} = r') / \Pr(\mathcal{D}_{j,k}) \\
=&\rho_{0,0}(s')\kappa_{0,0}(r')\phi_{0,0}(s',r')/\xi_{0,0}(0,0).
\end{aligned}$$

Now we show that the marginal likelihood ξ and ϕ indeed follow the recursive expression.

To this end, note that by definition, for $j = J$,

$$\phi_{j,k}(s,r) = \Pr(\mathcal{D}_{j,k}|S_{j,k} = s, R_{j,k} = r) = m(s,r).$$

Then also by definition for $j > 0$,

$$\begin{aligned}
\xi_{j,k}(s,r) &= \Pr(\mathcal{D}_{j,k}|S_{j-1,\lfloor k/2\rfloor} = s, R_{j-1,\lfloor k/2\rfloor} = r) \\
&= \sum_{s',r'} \Pr(S_{j,k} = s', R_{j,k} = r'|S_{j-1,\lfloor k/2\rfloor} = s, R_{j-1,\lfloor k/2\rfloor} = r) \cdot \Pr(\mathcal{D}_{j,k}|S_{j,k} = s, R_{j,k} = r) \\
&= \sum_{s',r'} \rho_{j,k}(s,s') \cdot \kappa_{j,k}(r,r')\phi_{j,k}(s',r')
\end{aligned}$$

and for $j = 0$,

$$\begin{aligned}
\xi_{j,k}(s,r) = \Pr(\mathcal{D}_{j,k}) &= \sum_{s',r'} \Pr(S_{j,k} = s', R_{j,k} = r') \cdot \Pr(\mathcal{D}_{j,k}|S_{j,k} = s, R_{j,k} = r) \\
&= \sum_{s',r'} \rho_{0,0}(s') \cdot \kappa_{0,0}(r')\phi_{0,0}(s',r')
\end{aligned}$$

Similarly, for $j < J$,

$$\begin{aligned}
\phi_{j,k}(s,r) &= \Pr(\mathcal{D}_{j,k}|S_{j,k} = s, R_{j,k} = r) \\
&= m(s,r) \cdot \Pr(\mathcal{D}_{j+1,2k}, \mathcal{D}_{j+1,2k+1}|S_{j,k} = s, R_{j,k} = r) \\
&= m_{j,k}(s,r) \cdot \xi_{j+1,2k}(s,r) \cdot \xi_{j+1,2k+1}(s,r).
\end{aligned}$$

This shows that the marginal posterior on the latent variables is a MT with states $\{0,1\} \times \{0,1\}$ the claimed transition matrix. The conditional posteriors of the regression coefficients and the errors follow directly from standard results on Bayesian linear regression with the NIG conjugate prior. □

S2. Posterior NIG-MG for multiple factors

With L factors, the design matrix is now

$$X = (\mathbb{1}_n, \boldsymbol{e}_2^{(1)}, \ldots, \boldsymbol{e}_{G_1}^{(1)}, \ldots, \boldsymbol{e}_2^{(L)}, \ldots, \boldsymbol{e}_{G_L}^{(L)}),$$

where $n = \sum_{l=1}^{L} \sum_{g=1}^{G_l} n_g$, $\mathbb{1}_n$ is a vector of n ones, and $\boldsymbol{e}_g^{(l)}$ is a binary vector where the hth element is equal to one if the hth observation belongs to the gth group for the lth factor, and equal to zero otherwise. We also define the following matrices for $(s, \boldsymbol{r}) = (s, r_1, r_2, \ldots, r_L) \in \{0,1\}^{L+1}$:

$$X(s, \boldsymbol{r}) = (s\mathbb{1}_n, r_1\boldsymbol{e}_2^{(1)}, \ldots, r_1\boldsymbol{e}_G^{(1)}, \ldots, r_L\boldsymbol{e}_2^{(L)}, \ldots, r_L\boldsymbol{e}_G^{(L)}),$$

$$\Lambda_j = \mathrm{diag}(1/\tau_j, \underbrace{1/\upsilon_j^{(1)}, \ldots, 1/\upsilon_j^{(1)}}_{G_1 - 1 \text{ copies}}, \ldots, \underbrace{1/\upsilon_j^{(L)}, 1/\upsilon_j^{(L)}, \ldots, 1/\upsilon_j^{(L)}}_{G_L - 1 \text{ copies}})$$

where $\upsilon_j^{(l)} = 2^{-\alpha j}\upsilon_l$ and

$$M(s, \boldsymbol{r}) = \begin{pmatrix} s \\ r_1\mathbb{1}_{G_1-1} \\ r_2\mathbb{1}_{G_2-1} \\ \vdots \\ r_L\mathbb{1}_{G_L-1} \end{pmatrix} \begin{pmatrix} s & r_1\mathbb{1}'_{G_1-1} & r_2\mathbb{1}'_{G_2-1} & \cdots & r_L\mathbb{1}'_{G_L-1} \end{pmatrix}.$$

The marginal likelihood for the node-specific model on (j,k) given $S_{j,k} = s$ and $\boldsymbol{R}_{j,k} = \boldsymbol{r}$ is

$$m_{j,k}(s, \boldsymbol{r}) = \frac{(\nu\sigma_0^2)^{\nu+1}\Gamma(\nu + n/2 + 1)}{(2\pi)^{n/2}\Gamma(\nu+1)} \cdot \frac{|\Lambda_j|^{1/2}}{|\Lambda_j^*(s, \boldsymbol{r})|^{1/2}} \cdot \left[\nu\sigma_0^2 + \Upsilon_{j,k}(s, \boldsymbol{r})\right]^{-\nu - n/2 - 1}. \tag{14}$$

where

$$\Upsilon_{j,k}(s, \boldsymbol{r}) = \left\{\boldsymbol{d}'_{j,k}\boldsymbol{d}_{j,k} - [\boldsymbol{\mu}_{j,k}^*(s, \boldsymbol{r})]'\Lambda_j^*(s, \boldsymbol{r})\boldsymbol{\mu}_{j,k}^*(s, \boldsymbol{r})\right\}/2,$$

$$\Lambda_j^*(s, \boldsymbol{r}) = X(s, \boldsymbol{r})'X(s, \boldsymbol{r}) + \Lambda_j, \quad \text{and} \quad \boldsymbol{\mu}_{j,k}^*(s, \boldsymbol{r}) = [\Lambda_j^*(s, \boldsymbol{r})]^{-1}[X(s, \boldsymbol{r})'\boldsymbol{d}_{j,k}].$$

Theorem 3. The joint posterior on $\{z_{j,k}, \boldsymbol{\beta}_{j,k}, S_{j,k}, \boldsymbol{R}_{j,k}, \sigma^2_{j,k} : (j,k) \in \mathcal{T}\}$ under the L-factor NIG-MG is as follows.

- The marginal posterior of the hidden states $\{(S_{j,k}, \boldsymbol{R}_{j,k}) : (j,k) \in \mathcal{T}\}$ is an MT defined on the product state-space $\{0,1\}^{L+1}$ with

  1. State transition probabilities:

$$\Pr(S_{j,k} = s', \boldsymbol{R}_{j,k} = \boldsymbol{r}' | S_{j-1,\lfloor k/2 \rfloor} = s, \boldsymbol{R}_{j-1,\lfloor k/2 \rfloor} = \boldsymbol{r}, \mathcal{D}) = \rho_{j,k}(s,s') \prod_{l=1}^{L} \kappa_{j,k}(r_l, r'_l) \cdot \frac{\phi_{j,k}(s', \boldsymbol{r}')}{\xi_{j,k}(s, \boldsymbol{r})},$$

     for $j = 1, \ldots, J$.

  2. Initial state probabilities:

$$\Pr(S_{0,0} = s, \boldsymbol{R}_{0,0} = \boldsymbol{r} | \mathcal{D}) = \rho_{0,0}(s) \prod_{l=1}^{L} \kappa_{0,0}(r_l) \cdot \frac{\phi_{0,0}(s, \boldsymbol{r})}{\xi_{0,0}(0, \boldsymbol{0})}.$$

- The conditional posterior of $\sigma^2_{j,k}$ given $S_{j,k}$ and $\boldsymbol{R}_{j,k}$ is:

$$[\sigma^2_{j,k} | S_{j,k}, \boldsymbol{R}_{j,k}, \mathcal{D}] \sim \text{Inv-Gamma}\left(\nu + 1 + \frac{n}{2}, \nu\sigma_0^2 + \Upsilon_{j,k}(S_{j,k}, \boldsymbol{R}_{j,k})\right).$$

- The posterior of $z_{j,k}, \beta^{(2)}_{1\,j,k}, \ldots, \beta^{(G_1)}_{1\,j,k}, \ldots, \beta^{(2)}_{L\,j,k}, \ldots, \beta^{(G_L)}_{L\,j,k}$ given $S_{j,k}$, $\boldsymbol{R}_{j,k}$ and $\sigma^2_{j,k}$ is given as follows

$$[z_{j,k}, \beta^{(2)}_{1\,j,k}, \ldots, \beta^{(G_1)}_{1\,j,k}, \ldots, \beta^{(2)}_{L\,j,k}, \ldots, \beta^{(G_L)}_{L\,j,k} \mid \sigma^2_{j,k}, S_{j,k}, \boldsymbol{R}_{j,k}, \mathcal{D}]$$
$$\sim \text{N}\left(\boldsymbol{\mu}^*_{j,k}(S_{j,k}, \boldsymbol{R}_{j,k}), \sigma^2_{j,k}\, M(S_{j,k}, \boldsymbol{R}_{j,k}) \circ [\Lambda^*_j(S_{j,k}, \boldsymbol{R}_{j,k})]^{-1}\right).$$

The mappings $\phi_{j,k}(s, \boldsymbol{r})$ and $\xi_{j,k}(s, \boldsymbol{r}) : \{0,1\}^{L+1} \to [0, +\infty)$ can be computed recursively through a pyramid algorithm as follows:

$$\phi_{j,k}(s, \boldsymbol{r}) = \begin{cases} m_{j,k}(s, \boldsymbol{r}) \cdot \xi_{j+1,2k}(s, \boldsymbol{r}) \cdot \xi_{j+1,2k+1}(s, \boldsymbol{r}) & \text{for } j = 0, 1, 2 \ldots, J-1 \\ m_{j,k}(s, \boldsymbol{r}) & \text{for } j = J, \end{cases}$$

$$\xi_{j,k}(s, \boldsymbol{r}) = \begin{cases} \sum_{s', \boldsymbol{r}'} \rho_{j,k}(s, s') \cdot \prod_{l=1}^{L} \kappa_{j,k}(r_l, r'_l) \cdot \phi_{j,k}(s', \boldsymbol{r}') & \text{for } j = 1, 2, \ldots, J \\ \sum_{s', \boldsymbol{r}'} \rho_{0,0}(s') \cdot \prod_{l=1}^{L} \kappa_{0,0}(r_l) \cdot \phi_{0,0}(s', \boldsymbol{r}') & \text{for } j = 0. \end{cases}$$

S3. Additional figures

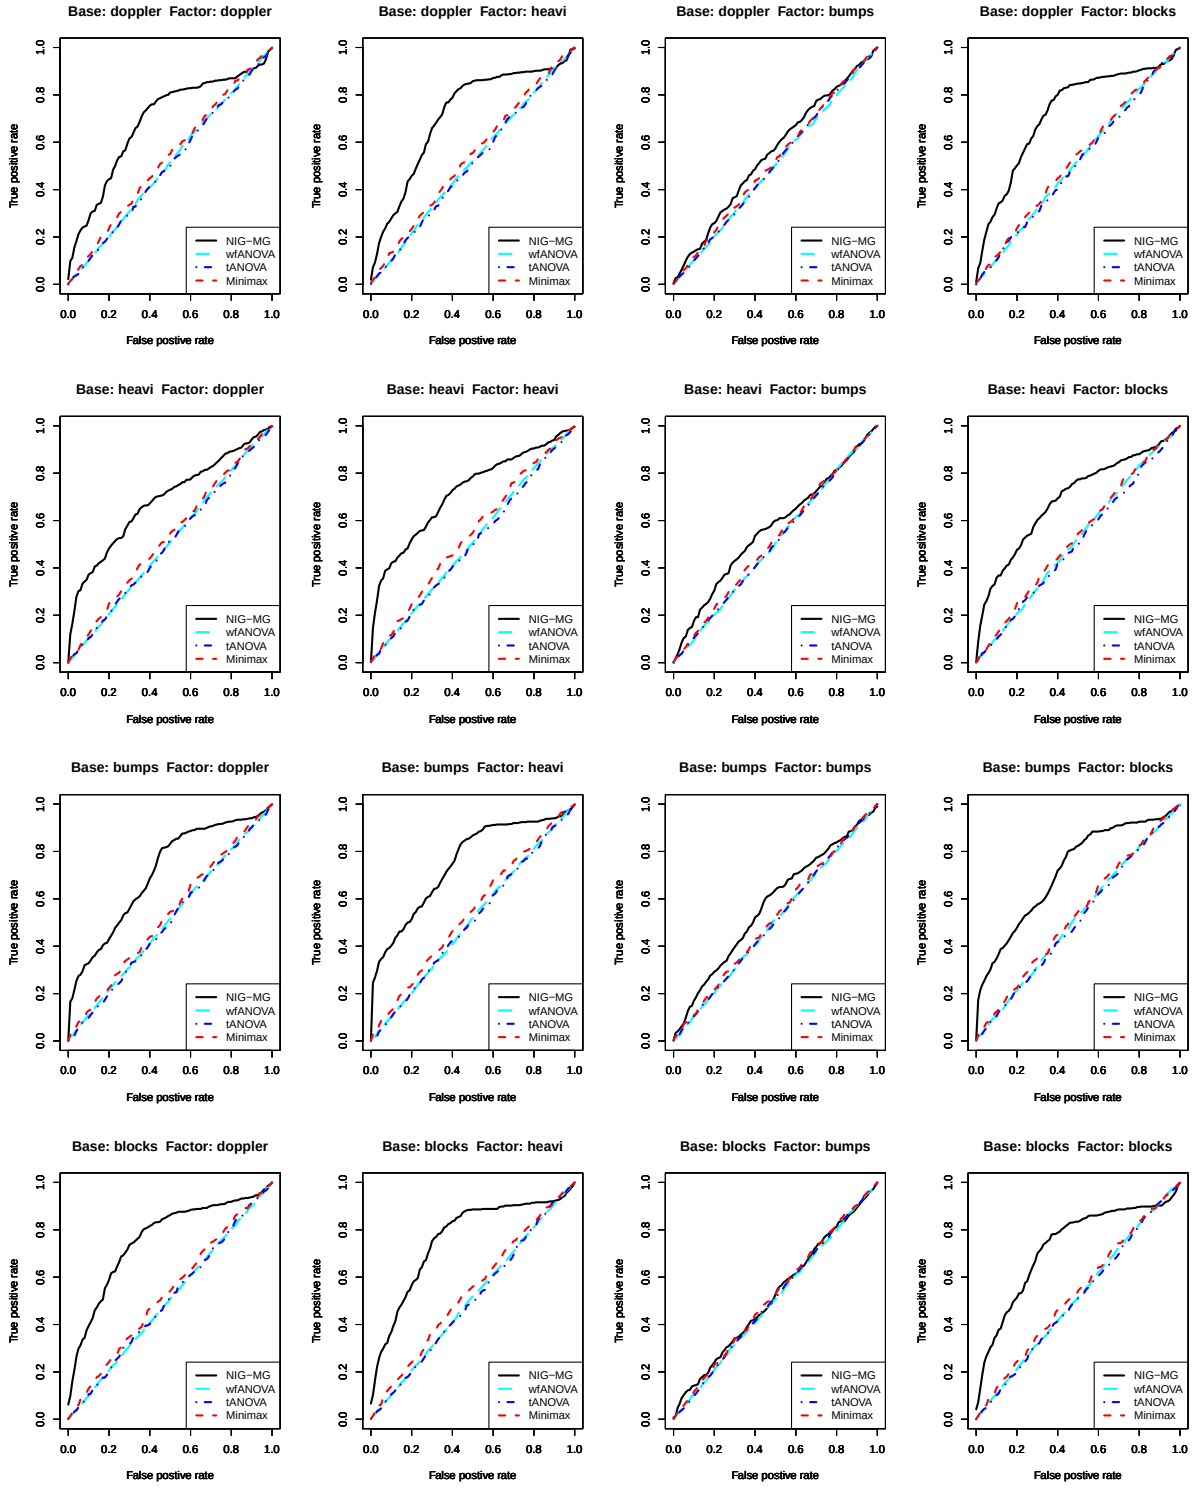

Figure S1: ROC curves (RSNR=0.5) for four test statistics when the baseline mean and the factor difference are all 16 combinations of the four signature functions.

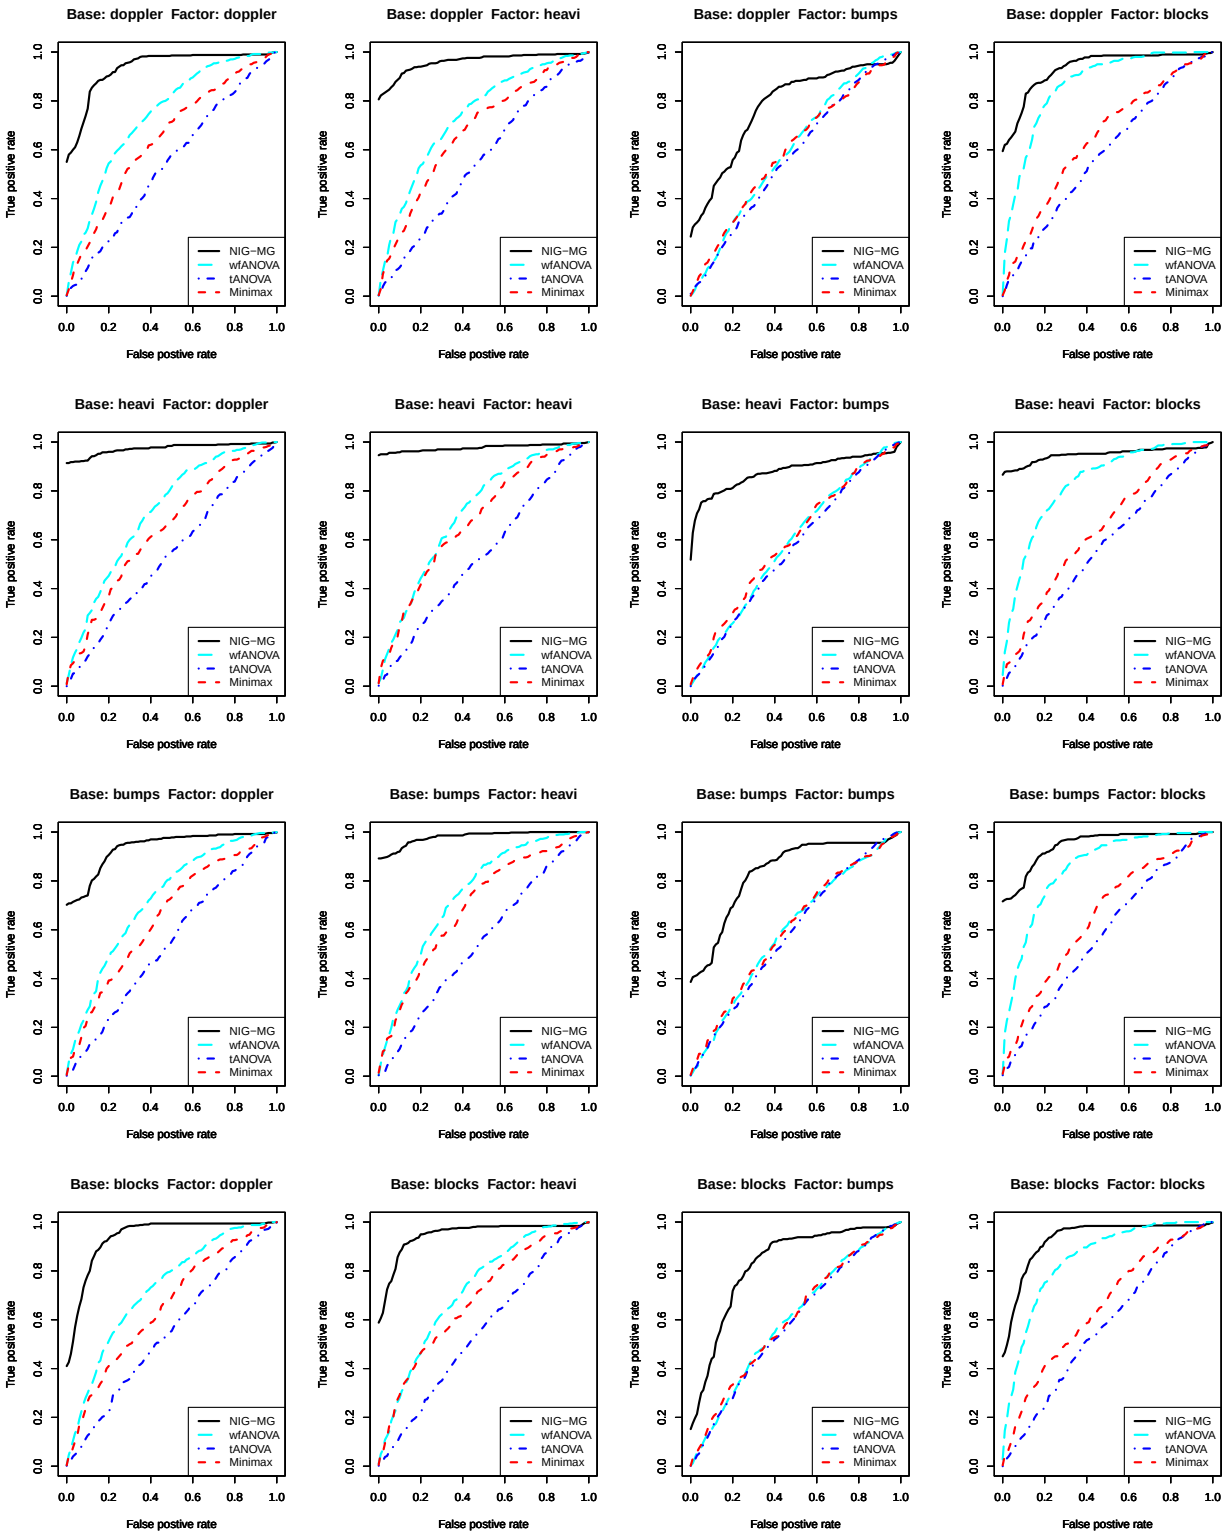

Figure S2: ROC curves (RSNR= 1.5) for four test statistics when the baseline mean and the factor difference are all 16 combinations of the four signature functions.